User Guide

hp StorageWorks Fabric OS 4.2.x Features

Second Edition (August 2004)

Part Number: AA-RV2HB-TE

This user guide provides procedures for setting up and configuring the optional features associated with Fabric OS v4.2.x, including Zoning, Advanced Performance Monitoring, ISL Trunking, Remote Switch and Extended Fabrics.



© Copyright 2004 Hewlett-Packard Development Company, L.P.

Hewlett-Packard Company makes no warranty of any kind with regard to this material, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. Hewlett-Packard shall not be liable for errors contained herein or for incidental or consequential damages in connection with the furnishing, performance, or use of this material.

This document contains proprietary information, which is protected by copyright. No part of this document may be photocopied, reproduced, or translated into another language without the prior written consent of Hewlett-Packard. The information contained in this document is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

Hewlett-Packard Company shall not be liable for technical or editorial errors or omissions contained herein. The information is provided "as is" without warranty of any kind and is subject to change without notice. The warranties for Hewlett-Packard Company products are set forth in the express limited warranty statements for such products. Nothing herein should be construed as constituting an additional warranty.

Fabric OS 4.2.x Features User Guide
Second Edition (August 2004)
Part Number: AA-RV2HB-TE

contents

About this Guide 7
Overview 8
Intended Audience 8
Related Documentation 8
Conventions 9
Document Conventions 9
Text Symbols 9
Equipment Symbols 10
Getting Help 12
HP Technical Support 12
HP Storage Web Site 12
HP Authorized Reseller 12

1 Fabric and Switch Management 13
Overview of Switch Management 14
Fabric Management Recommendations 15
High Availability 16
HA Features on the Core Switch 2/64 and SAN Director 2/128 16
SAN Switch 2/32 HA Features 17
SAN Switch 2/8V and SAN Switch 2/16V HA Features 18
User Access Levels 19
Fabric OS Command Line Interface 20
Advanced Web Tools 21
Fabric Manager 22
Fabric Watch 23
Fabric OS Access Layer (API) 24
Management Server 25

2 Advanced Zoning 27
Advanced Zoning Overview 28

Zone Definitions 30
Zone Objects 30
Zone Aliases 31
Zone Configurations 31
Zoning Schemes 32
Zone Enforcement 33
Software-Enforced Zoning 33
Hardware-Enforced Zoning 33
Best Practices for Hardware-Enforced Zoning 34
Hardware-Enforced Zoning Configurations 35
Rules for Configuring Zones 38
Fabric-Based Zoning 39
Zoning by Single HBA 40
Zoning by Application 40
Zoning by Operating System 40
Zoning by Port Allocation 41
No Fabric Zoning 41
The Effect of Zoning on RSCN Delivery 42
Using Zoning to Administer Security 43
Zone Management 44

3 Advanced Performance Monitoring 47
Overview of Advanced Performance Monitoring 48
Types of Performance Monitoring 49
AL_PA Monitoring 49
End-to-End Monitoring 49
Where to Add End-to-End Monitors 51
Filter-Based Monitoring 53
Standard Filter-Based Monitors 53
User-Defined Filter-Based Monitors 53
Operating Advanced Performance Monitoring with Advanced Web Tools 54
Predefined Performance Graphs 55
User-Defined Graphs 56
Performance Monitoring with Telnet Commands 58

4 ISL Trunking 61
Overview of ISL Trunking 62
ISL Trunking Architecture 63
Traffic Routing 64

Trunking Groups 64
Trunking ISLs 64
Trunking Ports 65
Setting Up ISL Trunking in a Fabric 71
Evaluating Data Traffic Patterns 71
Designing the Fabric to Optimize Use of ISL Trunking 72
Managing ISL Trunking 74
Commands Related to ISL Trunking 76
Frequently Asked Questions About ISL Trunking 77

5 Extended Fabrics 79

6 Remote Switch 81

Glossary 83

Index 103

Figures

1 Fabric with Three Zones 29
2 Nonoverlapping Hardware Enforced Zoning 35
3 Overlapping Hardware-Enforced Zoning 36
4 Zoning with Hardware Assist (Mixed Port and WWN Zones) 37
5 Overlapping Hardware Enforced Zoning with Soft Porting 37
6 Setting End-to-End Monitors on a Port 51
7 Proper Placement of End-to-End Performance Monitors 52
8 Accessing Performance Graphs 56
9 Eight Performance Monitoring Graphs 57
10 Distribution of Traffic Over an ISL Trunking Group 63
11 Port Groupings on a SAN Switch 2/8 EL 66
12 Quad and Port Groupings on a SAN Switch 2/8V 66
13 Quad and Port Groupings on a SAN Switch 2/16 67
14 Quad and Port Groupings on a SAN Switch 2/16V 67
15 Quad and Port Groupings on a SAN Switch 2/32 68
16 Port Groupings on the Port Card on a Core Switch 2/64 69
17 Port Groupings on the Port Card on a SAN Director 2/128 70
18 Trunk Information Tab in the Switch Admin Window of Advanced Web Tools 75

Tables

1 Document Conventions 9
2 Maximum Allowed Sessions 19
3 RSCN Delivery 42
4 Commands Used for Managing Zoning 45
5 Predefined Performance Graphs 55
6 APM Telnet Commands. 58
7 Fabric OS Commands Related to ISL Trunking 76
8 Frequently Asked Questions about ISL Trunking 77

about this guide

This user guide provides guidelines for using HP StorageWorks Fabric OS 4.2.x software options.“About this Guide” topics include:

- Overview, page 8
- Conventions, page 9
- Getting Help, page 12

Overview

This section covers the following topics:

- Intended Audience
- Related Documentation

Intended Audience

This book is intended for use by system administrators and technicians who are experienced with the following:

- HP StorageWorks Fibre Channel SAN switches
- Fabric Operating System version 4.2.x or later

Related Documentation

Documentation, including white papers and best practices documents, is available via the HP website. Please go to:

http://www.hp.com/country/us/eng/prodserv/storage.html

To access v4.2.x related documents:

1. Locate the "Networked storage" section of the web page.
2. Under "Networked storage," go to the "By type" subsection.
3. Click **SAN infrastructure**. The SAN infrastructure page displays.
4. Locate the Fibre Channel Switches section.
5. Locate the B-Series Fabric subsection, then go to the "Enterprise Class" subsection.
6. To access version 4.2.x documents (like this document), select **SAN Director 2/128** or **SAN Director 2/128 PowerPak**. The switch overview page displays.
7. Go to the "product information section," located on the far right side of the web page.
8. Click **technical documents**.
9. Follow the onscreen instructions to download the applicable documents.

Conventions

Conventions consist of the following:

- Document Conventions
- Text Symbols
- Equipment Symbols

Document Conventions

This document follows the conventions in Table 1.

Table 1: Document Conventions

Convention	Element
Blue text: Figure 1	Cross-reference links
Bold	Menu items, buttons, and key, tab, and box names
Italics	Text emphasis and document titles in body text
`Monospace font`	User input, commands, code, file and directory names, and system responses (output and messages)
`Monospace, italic font`	Command-line and code variables
Blue underlined sans serif font text (http://www.hp.com)	Web site addresses

Text Symbols

The following symbols may be found in the text of this guide. They have the following meanings:

WARNING: Text set off in this manner indicates that failure to follow directions in the warning could result in bodily harm or death.

Caution: Text set off in this manner indicates that failure to follow directions could result in damage to equipment or data.

Tip: Text in a tip provides additional help to readers by providing nonessential or optional techniques, procedures, or shortcuts.

Note: Text set off in this manner presents commentary, sidelights, or interesting points of information.

Equipment Symbols

The following equipment symbols may be found on hardware for which this guide pertains. They have the following meanings:



Any enclosed surface or area of the equipment marked with these symbols indicates the presence of electrical shock hazards. Enclosed area contains no operator serviceable parts.

WARNING: To reduce the risk of personal injury from electrical shock hazards, do not open this enclosure.



Any RJ-45 receptacle marked with these symbols indicates a network interface connection.

WARNING: To reduce the risk of electrical shock, fire, or damage to the equipment, do not plug telephone or telecommunications connectors into this receptacle.



Any surface or area of the equipment marked with these symbols indicates the presence of a hot surface or hot component. Contact with this surface could result in injury.

WARNING: To reduce the risk of personal injury from a hot component, allow the surface to cool before touching.



Power supplies or systems marked with these symbols indicate the presence of multiple sources of power.

WARNING: To reduce the risk of personal injury from electrical shock, remove all power cords to completely disconnect power from the power supplies and systems.



Any product or assembly marked with these symbols indicates that the component exceeds the recommended weight for one individual to handle safely.

WARNING: To reduce the risk of personal injury or damage to the equipment, observe local occupational health and safety requirements and guidelines for manually handling material.

Getting Help

If you still have a question after reading this guide, contact an HP authorized service provider or access our web site at http://www.hp.com.

HP Technical Support

Telephone numbers for worldwide technical support are listed on the following HP web site: http://www.hp.com/support/. From this web site, select the country of origin.

Note: For continuous quality improvement, calls may be recorded or monitored.

Be sure to have the following information available before calling:

- Technical support registration number (if applicable)
- Product serial numbers
- Product model names and numbers
- Applicable error messages
- Operating system type and revision level
- Detailed, specific questions

HP Storage Web Site

The HP web site has the latest information on this product, as well as the latest drivers. Access storage at: http://www.hp.com/country/us/eng/prodserv/storage.html. From this web site, select the appropriate product or solution.

HP Authorized Reseller

For the name of your nearest HP authorized reseller:

- In the United States, call 1-800-345-1518
- In Canada, call 1-800-263-5868
- Elsewhere, see the HP web site for locations and telephone numbers: http://www.hp.com.

Fabric and Switch Management

1

This chapter explains the different methods used to manage an HP StorageWorks switch, and provides the following information:

- Overview of Switch Management, page 14
- Fabric Management Recommendations, page 15
- High Availability, page 16
- User Access Levels, page 19
- Fabric OS Command Line Interface, page 20
- Advanced Web Tools, page 21
- Fabric Manager, page 22
- Fabric Manager, page 22
- Fabric OS Access Layer (API), page 24
- Management Server, page 25

Overview of Switch Management

HP StorageWorks switches can be managed using several local and remote access methods. To manage a switch, you must have access to one of the available management methods:

- Fabric OS Command Line Interface (CLI)
- Advanced Web Tools
- Fabric Manager
- Fabric Watch
- Fabric OS Access Layer (API)
- Management Server

Telnet, SNMP, and Advanced Web Tools require that the switch be accessible using a network connection. The network connection can be from the switch Ethernet port (out of band) or from the fibre channel (in band). The switch must be configured with an IP address to allow for the network connection. Refer to the installation guide for your specific switch for information on physically connecting to the switch.

Fabric Management Recommendations

HP makes the following recommendations for managing switches and fabrics:

- Mixed fabrics

 In a mixed fabric containing Fabric OS v4.x, v3.x, and v2.x switches, manage the fabric via the switch with the latest Fabric OS version as the first criteria, then the most advanced as the second criteria. For example, use the v4.2.x switches running on the SAN Director 2/128 as the primary FCS (when running Secure Fabric OS), the location to perform zoning tasks, and the time server.

- Multiple connections

 Switches can be accessed simultaneously from different connections (for example, Advanced Web Tools, CLI, and API). If this happens, changes from one connection might not be updated to the other, and some modifications might be lost. When accessing the switch via multiple connections, make sure that you do not overwrite changes.

- Fabric-level tasks

 A number of management tasks are designed to make fabric-level changes. When executing fabric-level configuration tasks (like issuing Zoning commands) make sure to allow time for the changes to propagate across the fabric before executing any sequential tasks. For example in a large fabric, this might take up to a few minutes.

- Command lettercasing

 In this guide, all commands are shown as lowercase. This is because Fabric OS v4.2.x, unlike previous versions of the Fabric OS, is case-sensitive. For backward compatibility, you can enter the commands using the legacy mixed-case notation (for example, `setGbicMode`). Lowercase is recommended and always works on both v3.x and v4.x.

- High availability failover

 Prior to initiating `hafailover`, make sure that any previously executed commands have completed. It might be necessary to reissue the command if it has failed, wholly or partially, due to the failover.

High Availability

This section provides information on the HP StorageWorks High Availability (HA) feature.

HA Features on the Core Switch 2/64 and SAN Director 2/128

HA features include:

- Redundant AC cables
- Redundant, hot-swappable power supplies
- Redundant, hot-swappable fans
- Hot-swappable port blades
- Hot-swappable Control Processor (CP) blades
- Hot-swappable WWN card
- Nondisruptive failover of CP blades
- Background health monitor and automatic failover in case of active CP failure
- Nondisruptive firmware download (external server to flash)
- Nondisruptive code activation

A key part of the HA feature is a software component known as the HA Manager. Via the HA Manager, the standby CP maintains the same state settings in real-time as the active CP. The moment the standby CP needs to take over operation of a switch, it is essentially in the same state as the active CP and can seamlessly assume control of the switch. This enables nondisruptive failover and code activation on the Core Switch 2/64 and SAN Director 128 models. As a result, there is no effect whatsoever on the flow of data between logged-in hosts and storage devices. Read and write tasks proceed with no delay or interruption. No frames are dropped. No devices that are currently logged in have to re-login.

The HA Manager uses a “take-away” model for changing control of the active CP to the standby CP. The active CP is not relied upon to discover that something is wrong and “give-away” control; rather, the HA Manager on the standby CP includes mechanisms to detect a failure in the active CP. This enables the switch to seamlessly continue processing (on the standby CP) in the event of active CP failure.

During a failover event, there is some delay in response to requests for fabric services, such as queries to the Name Server, bringing up new links (whether ISLs or edge devices), or processing fabric login (FLOGI) requests. This delay is in the

range of only 5 to 10 seconds, substantially less than the limit imposed by Fibre Channel standards, which indicate that the switch must respond in less than twice the resource allocation time out value (R_A_TOV). This type of delay primarily affects those hosts whose state is in the process of changing: mounting a device, for example.

SAN Switch 2/32 HA Features

The SAN Switch 2/32 delivers a number of HA features, including:

- Redundant AC cables
- Redundant, hot-swappable power supplies
- Redundant, hot-swappable fans
- Nondisruptive firmware download (external server to flash)
- Nondisruptive code activation

The same mechanisms used to provide nondisruptive failover on the Core Switch 2/64 and SAN Director 128 models are used to perform nondisruptive reboots and code activation on the SAN Switch 2/32 model.

The service running on the SAN Switch 2/32 is a normal active CP instance image; however, since there is no standby CP, state updates are not maintained.

When a reboot or code activation process is initiated, a reboot management utility is launched. This program creates a standby image of the current switch state. The reboot management program then initiates a synchronization, causing all components of the active image to replicate to the standby image. As the state updates are synchronized, messages are saved in Flash.

When all components have finished synchronizing, the active image is disengaged and the actual reboot or code activation operation occurs. Control of the switch is passed seamlessly to the standby image. As a result, there is no effect whatsoever on the flow of data between logged in hosts and storage devices. Read and write tasks proceed with no delay or interruption whatsoever. No frames are dropped. No devices that are currently logged in have to re-login.

SAN Switch 2/8V and SAN Switch 2/16V HA Features

The SAN Switch 2/8V and SAN Switch 2/16V models deliver a number of HA features, including:

- Redundant AC cables
- Nondisruptive firmware download (external server to flash)
- Nondisruptive code activation

The same mechanisms used to provide nondisruptive failover on the Core Switch 2/64 and SAN Director 128 models are used to perform nondisruptive reboots and code activation on the SAN Switch 2/8V and SAN Switch 2/16V models.

The service running on the SAN Switch 2/8V and SAN Switch 2/16V is a normal active CP instance image; however, since there is no standby CP, state updates are not maintained.

When a reboot or code activation process is initiated, a reboot management utility is launched. This program creates a standby image of the current switch state. The reboot management program then initiates a synchronization, causing all components of the active image to replicate to the standby image. As the state updates are synchronized, messages are saved in Flash.

When all components have finished the synchronization, the active image is disengaged and the actual reboot or code activation operation occurs. Control of the switch is passed seamlessly to the standby image. As a result, there is no effect whatsoever on the flow of data between logged in hosts and storage devices. Read and write tasks proceed with no delay or interruption whatsoever. No frames are dropped. No devices that are currently logged in have to re-login.

User Access Levels

There are four levels of user access to HP StorageWorks switches:

- root
- factory
- admin
- user

Not all commands are available to all levels. Commands are assigned a minimum login level to execute. Root level has access to all commands; user level has limited access to commands.

Note: Because the admin account provides access to all the commands needed to manage and configure a switch or fabric, it is the recommended login level.

In Fabric OS v3.x and earlier, multiple user access to a switch is limited. Each switch enables only a single session per management access method, regardless of user level. Switches can, however, be accessed simultaneously from different connections (for example, through the CLI and Advanced Web Tools). If this happens, changes from one connection might not be updated to the other, and some changes might be lost. When accessing the switch simultaneously via multiple connections, make sure that you do not overwrite the work of another connection.

See Table 2 for the maximum number of simultaneous sessions allowed.

Table 2: Maximum Allowed Sessions

User Name	Maximum Number of Simultaneous Sessions
root	4
factory	4
admin	2
user	4

Fabric OS Command Line Interface

The Fabric OS command line interface (CLI) accessed through telnet or serial console provides the user with the full range of management capability on an HP StorageWorks switch. The CLI enables an administrator to monitor and manage entire fabrics, individual switches, and ports from a standard workstation. The entire suite of Fabric OS features and capabilities is available across an entire fabric, from a single access point.

Access is controlled by a switch-level password for each user level (root, factory, admin, user). The commands available through the CLI are based on the user's login level and the license keys used to unlock certain features.

All configuration and management tasks are available using the admin or user level; the root and factory levels should be used with caution. The *HP StorageWorks Fabric OS 4.2.x Command Reference Guide* lists all the commands available to the user- and admin-level users.

The Fabric OS CLI is the complete fabric management tool for HP StorageWorks SANs, providing the following advantages:

- Access to the full range of Fabric OS features, based on which license keys you purchase
- A full set of tools for configuring, monitoring, dynamic provisioning, and daily management of every aspect of storage area networks.
- The ability to configure and manage the fabric on multiple efficient levels
- The ability to identify, isolate, and manage SAN events across every switch in the fabric
- The ability to manage switch licenses
- The ability to perform Fabric Stamping

For more information about using the CLI, refer to the *HP StorageWorks Fabric OS 4.2.x Command Reference Guide* and the *HP StorageWorks Fabric OS 4.2.x Procedures User Guide.*

Advanced Web Tools

Advanced Web Tools provides a Graphical User Interface (GUI) that enables you to monitor and manage individual switches and ports from a standard workstation. It is a licensed product that runs on Fabric OS. All switches in the fabric are displayed in the main window of Advanced Web Tools, including switches that do not have an Advanced Web Tools license; however, only those switches that have an Advanced Web Tools license installed can be managed through Advanced Web Tools (other switches must be managed through telnet).

Following are some of the features that make Advanced Web Tools an important part of the switch management:

- Advanced Web Tools can be used simultaneously with Fabric OS CLI commands. Simply open a second window and you can take advantage of the benefits of both interfaces at the same time.
- Advanced Web Tools can help you find the appropriate Fabric OS CLI command to perform a specific function. For example, you can perform a function using Advanced Web Tools, and watch in a second window as the Fabric OS CLI commands are displayed.
- Advanced Web Tools can be used from a standard workstation, allowing you to be virtually "in front of" any fabric, switch, or port.
- Advanced Web Tools allows you to zone your configuration by dragging addresses and port numbers rather than laboriously typing them out.
- Advanced Web Tools provides the Performance Monitor feature, enabling you to view the status and traffic of a switch or port in seconds via a variety of effective graphs. Refer to "Advanced Performance Monitoring" on page 47 for more information on this feature.
- Advanced Web Tools is easy and intuitive to use.

For more information, refer to the *HP StorageWorks Advanced Web Tools 4.2.x User Guide*.

Fabric Manager

Like Advanced Web Tools, Fabric Manager provides a GUI for monitoring and managing switches and ports from a standard workstation; however, in addition to individual switches and ports, Fabric Manager also enables you to monitor and manage entire fabrics.

Fabric Manager can be used to manage fabrics containing integrated fabrics in addition to individual HP StorageWorks switches. Fabric Manager provides high-level information about all switches in the fabric.

All switches in the fabric are represented in the main window of Fabric Manager, but only those with an Advanced Web Tools license can be managed through Fabric Manager.

Fabric Manager is the complete SAN management tool for HP StorageWorks SANs, providing the following advantages to administrators:

- A highly scalable, Java-based application that manages multiple switches and multiple fabrics in real-time
- Time savings by enabling global integration and execution of processes across multiple fabrics, through a single-point SAN management platform
- More effective management by providing rapid access to critical SAN information across both Fabric OS SANs and enhanced Fabric OS SANs
- The ability to identify, isolate, and manage SAN events across multiple switches and fabrics
- Drill-down capability to individual SAN components, through tightly coupled Advanced Web Tools and Fabric Watch integration
- Discovery of all SAN components and ability to view the real-time state of all HP StorageWorks fabrics
- Multifabric administration of Secure Fabric OS SANs through a single encrypted console
- The ability to monitor interswitch links (ISLs)
- The ability to manage switch licenses
- The ability to perform Fabric Stamping

For more information, refer to the *HP StorageWorks Fabric Manager 4.1.1x User Guide* and integrated online help.

Fabric Watch

Fabric Watch is a licensed product that monitors the performance and status of Fibre Channel networks and HP StorageWorks switches, alerting you when problems arise before they become costly failures. SAN managers can configure Fabric Watch software to monitor any of the following:

- Fabric events (such as topology reconfigurations and zone changes)
- Physical switch conditions (such as fans, power supplies, and temperature)
- Port behavior (such as state changes, errors, and performance)
- SFP behavior

With Fabric Watch software, you can place limits, or *thresholds*, on the behavior of different switch and fabric elements. Fabric Watch then monitors these behavior variables, or *counters*, and issues an alarm to address problems when a counter exceeds a threshold. An alarm can email you or lock out a port log, depending on how you configure the alarm.

For more information, refer to the *HP StorageWorks Fabric Watch 4.2.x User Guide*.

Fabric OS Access Layer (API)

The Fabric OS Access Layer is an Application Programming Interface (API) that enables any application to access critical information about an HP StorageWorks SAN. With Fabric Access, an application can query or control individual switches or the entire fabric.

The Fabric OS Access Layer has the following advantages:

- It uses third-party software to manage an HP fabric
- It can create SAN management applications specific to your needs, using the tools available in the Fabric OS.

Management Server

The Fabric OS includes a Distributed Management Server. The Management Server (MS) allows a SAN management application to retrieve information and administer the fabric and interconnected elements, such as switches, servers, and storage devices. The MS is located at the Fibre Channel well-known address *FFFFFAh*.

The implementation of the Management Server provides the following management services:

- Fabric Configuration Service

 Provides basic configuration management for topology information (referred to as Topology Discovery).

- Unzoned Name Server access

 Provides a management view of the Name Server information for all devices in a fabric, regardless of the active zone set.

- Fabric Zone Service
- FDMI

The services provided by the MS assist in the autodiscovery of switch-based fabrics and their associated topology. A client of the MS can determine basic information regarding the switches in the fabric and use this information to construct topology relationships.

In addition, the basic configuration services provided by the management server allow certain attributes associated with switches to be obtained and, in some cases, modified. For example, logical names identifying switches can be registered with the Management Server.

For more information, refer to the *HP StorageWorks Fabric OS 4.2.x Procedures User Guide*.

Note: The Management Server is disabled if the switch is in secure mode. Refer to the *HP StorageWorks Secure Fabric OS 4.2.x User Guide* for more information.

Advanced Zoning

2

This chapter contains the following information:

- Advanced Zoning Overview, page 28
- Zone Definitions, page 30
- Zone Enforcement, page 33
- Rules for Configuring Zones, page 38
- Fabric-Based Zoning, page 39
- The Effect of Zoning on RSCN Delivery, page 42
- Using Zoning to Administer Security, page 43
- Zone Management, page 44

Advanced Zoning Overview

Advanced Zoning is a licensed HP product that allows partitioning of a SAN into logical groupings of devices that access each other. These logical groupings are called *zones*.

You can use Advanced Zoning to customize environments and optimize resources:

- Customize Environments—You can use zones to create logical subsets of the fabric to accommodate environments such as closed user groups or functional areas within the fabric. For example, you can identify selected devices within a zone for the exclusive use of zone members, or you can define a zone to create separate test or maintenance areas within the fabric.
- Optimize Resources—You can use zones to logically consolidate equipment for efficiency or to facilitate time-sensitive functions: for example, to create a temporary zone to back up nonmember devices.

A zone is a specified group of fabric-connected devices, also called *zone objects*, that have access to one another. Zone objects are grouped into zones, and zones are grouped into a zone configuration. Zones can overlap; that is, a zone object can belong to more than one zone and a fabric can have multiple zones. A switch can have any number of resident zone configurations; however, only one configuration can be enabled at a time.

Figure 1 illustrates a fabric with three zones.

Any zone object connected to the fabric can be included in one or more zones. For example, in Figure 1, the RAID device is included in the Blue zone and the Green zone.

Objects in a zone can communicate only with other objects in the same zone. For example, in Figure 1, Server1 can communicate only with the Loop2 devices; Server3 can communicate only with the RAID and the Loop2 devices. Note that Server3 can communicate with the RAID device even though it is in the Blue zone (which is inaccessible to Server3), because the RAID device is also in the Green zone (which is the same zone as Server3).

Note: After zoning is enabled, if a device is not explicitly defined in a zone, that device is isolated and inaccessible by other devices in the fabric.

For example, in Figure 1, the Loop1 JBODs are not assigned to a zone; thus, they are not active in the fabric if zoning is enabled, and no other device can access them.

Zones can be configured dynamically and can vary in size, depending on the number of fabric-connected devices.

Zoning can be disabled at any time. When zoning is disabled, the fabric is in nonzoning mode, and devices can freely access other devices in the fabric.

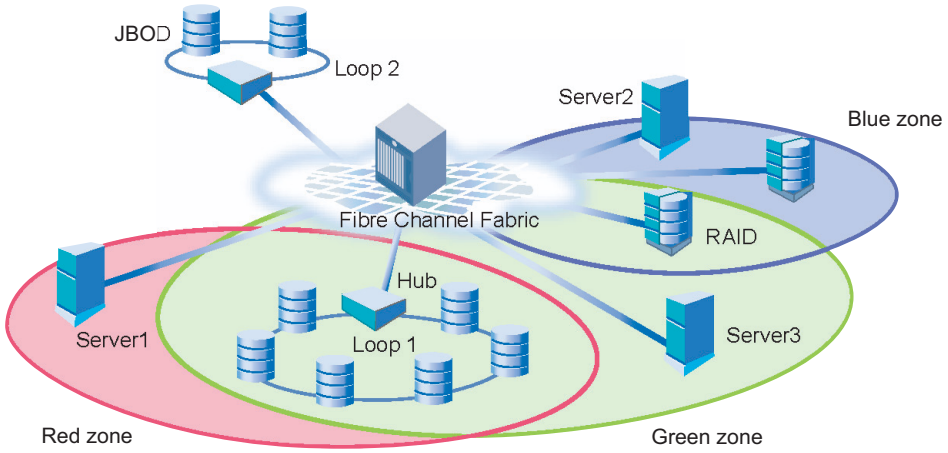


Figure 1: Fabric with Three Zones

All devices connected to a fabric can be configured into one or more zones. Every zone has at least one zone object. Empty zones are not allowed.

The memory constraints for zoning are far larger than the number of devices that can be connected to a fabric. The zone size for Fabric OS v4.2.x is 127726 bytes. The zone size for Fabric OS v3.x and v2.x is 98232 Bytes.

Zone configurations are consistent across reboots and power cycles. If two switches are connected in a fabric, they can become isolated (for example, due to an ISL failure); however, when rejoined, they maintain the same fabric configuration unless one of the switches has had a configuration change.

Zone Definitions

This section describes the following zoning terminology:

- zone objects
- zone aliases
- zone configurations
- zoning schemes

Zone Objects

As stated earlier, a *zone object* is any device in a zone:

- Physical port number (area number) on the switch
- Node World Wide Name (WWN)
- Port World Wide Name

Zone objects identified by port number or area number are specified as a pair of decimal numbers, *d,area*, where *d* is the domain ID of the switch and *area* is the area number on that switch. For example, in Fabric OS v4.1.x, “3, 46” specifies port 14 on slot number 3 or area 46. In Fabric OS v3.1.x, “3,13” specifies port 13 on switch number 3.

When the physical port number specifies a zone object, then all devices connected to that port are in the zone. If this physical port is an arbitrated loop, then all devices on the loop are part of the zone.

World Wide Names are specified as eight-digit hexadecimal numbers, separated by colons: for example, “10:00:00:90:69:00:00:8a”. When a node name specifies a zone object, all ports on that device are in the zone. When a port name specifies a zone object, only that single port is in the zone.

The types of zone objects used to define a zone can be mixed and matched. For example, a zone defined with the zone objects “2,12; 2,14; 10:00:00:80:33:3f:aa:11” contains the devices connected to switch 2, ports 12 and 14, and a device with the WWN (either node name or port name) “10:00:00:80:33:3f:aa:11” that is connected on the fabric.

Zone Aliases

A *zone alias* is a name assigned to a device or group of devices. By creating an alias, you can assign a familiar name to a device or group multiple devices into a single name. This can simplify cumbersome data entry and allows an intuitive naming structure (such as using "NT_Hosts" to define all NT hosts in the fabric).

Zone aliases also simplify repetitive entry of zone objects such as port numbers or a WWN. For example, the name "Eng" could be used as an alias for "10:00:00:80:33:3f:aa:11".

A useful convention is to name zones for the initiator they contain. For example, if you use the alias SRV_MAILSERVER_SLT5 to designate a server, host name mailserver, in PCI slot 5, then the alias for the associated zone would be ZNE_MAILSERVER_SLT5. This clearly identifies the server HBA that is associated with the zone.

Zone configuration naming is more flexible. In general, one configuration should be named PROD_*fabricname*, where *fabricname* is the name that the fabric has been designated. The purpose of the PROD configuration is to easily identify the configuration that can be implemented and provide the most generic services. If other configurations are used for specialized purposes, names such as "BACKUP_A," "RECOVERY_2," and "TEST_18jun02" can be used.

Zone Configurations

A *zone configuration* is a group of one or more zones. A zone can be included in more than one zone configuration. When a zone configuration is in effect, all zones that are members of that configuration are in effect.

The following terms are used to differentiate between types of zone configurations:

- Defined Configuration—The complete set of all zone objects that have been defined in the fabric.
- Effective Configuration—A single zone configuration that is currently in effect. The effective configuration is built when an administrator enables a specified zone configuration.
- Saved Configuration—A copy of the defined configuration plus the name of the effective configuration, which is saved in flash memory by the `cfgSave` command. There might be differences between the saved configuration and the defined configuration if the system administrator has modified any of the zone definitions and has not saved the configuration.

- Disabled Configuration—The effective configuration is removed from flash memory.

On power-up, the switch automatically reloads the saved configuration, and if a configuration was in effect when it was saved, the same configuration is reinstated with an autorun of the `cfgEnable` command.

Zoning Schemes

Zoning schemes define the method used to establish a zone. A zone is established by identifying zone objects by using one or more of the following schemes:

- Domain/port level—All members of this zone are specified by domain ID, port number, or area number, or aliases of *domain, area*.
- World Wide Name (WWN) level—All members of this zone are specified only by World Wide Name (WWNs) or aliases of WWNs. They can be node or port versions of the WWN.
- Mixed zoning—A zone containing members specified by a combination of domain/port (or area) number and WWN.

Zone Enforcement

Depending on the zoning scheme, zones are either software enforced or hardware enforced. When a zone is specified by either port number (area number) or WWN, but not both it is a hardware enforced zone.

HP Zoning software limits access to data by segmenting a fabric into virtual private SANs. Available on both 1 Gbit/sec and 2 Gbit/sec platforms, software-enforced zoning prevents hosts from discovering unauthorized target devices, while hardware-enforced zoning prevents a host from ever accessing a device it is not authorized to access.

Software-Enforced Zoning

Software-enforced zoning employs the name server to not return any information to an unauthorized initiator in response to a name server query. This is referred to as *software-enforced zoning* (or *soft zoning*, *name server zoning*, or *fabric-based zoning*). Software-enforced zoning is always active whenever a zone configuration is in effect.

When an initiator queries the name server for accessible devices in the fabric, the name server returns only those devices that are in the same zone as the initiator. Devices that are not part of the zone are not returned as accessible devices.

It is important to note that software-enforced zoning does not prohibit access to the device. If an initiator has knowledge of the network address of a target device, it does not need to query the name server to access it, which could lead to undesired access to a target device by unauthorized hosts.

Software-enforced zoning is exclusively enforced through selective information presented to end nodes through the fabric Simple Name Server (SNS).

Hardware-Enforced Zoning

HP StorageWorks switches augment Name Server zoning with hardware enforcement. The exact methodology varies on different switch models.

Hardware-enforced zoning (or *hard zoning*) means that each frame is checked by hardware before it is delivered to a zone member and discarded if there is a zone mismatch.

Hardware-enforced zoning is enforced at the ASIC level. Each ASIC maintains a list of source port IDs that have permission to access any of the ports on that ASIC.

When hardware-enforced zoning is active, the switch monitors the communications and blocks any frames that do not comply with the effective zone configuration. This blocking is performed at the transmit side of the port where the destination device is located. A zone configuration typically comprises many zones, and the hardware enforcement decision is made on a per-zone basis. The use of hardware enforcement is decided by the Fabric OS whenever the fabric membership or zone configuration changes.

The switch models and hardware zoning methodology for each are:

- HP StorageWorks 1 GB switches—Enable hardware-enforced zoning only on domain/port zones; WWN or mixed zones are not hardware enforced. Any domain/port zone that overlaps a mixed or WWN zone is not hardware enforced.

 Note that an overlap occurs when a member specified by WWN is connected to a port in a domain/port zone. The domain/port zone loses its hardware enforcement even though a review of the zone configuration would not indicate this.

- HP StorageWorks 2 GB switches (includes Core Switch 2/64 and SAN Director 2/128 models)—Enable hardware-enforced zoning on domain/port zones and WWN zones. Overlap of like zone types does not result in the loss of hardware enforcement. Overlap with another zone type results in the loss of hardware enforcement.

 As in the HP StorageWorks 1 GB switches, connecting a device specified by WWN into a port specified in a domain/port zone results in loss of the hardware enforcement in both zones.

Note: In mixed fabrics, each switch type continues to enable hardware enforcement of zoning according to the conditions established above. Use `portzoneshow` to learn which switch type is attached.

Best Practices for Hardware-Enforced Zoning

Sometimes determining the applicability of hardware enforcement can be difficult. Three basic cases have simple rules for assuring that hardware enforcement is used:

- For fabrics containing HP StorageWorks 1 GB switches, always use domain/port identifiers; never identify a zone member by WWN.

- In fabrics containing HP StorageWorks 2 GB switches (includes Core Switch 2/64 and SAN Director 2/128 models) use either WWN or domain/port zones.
- For mixed fabrics, use domain/port identifiers.

Note: If you must use WWN identifiers, always place disk and tape units (targets) on HP StorageWorks 2 GB switches (including Core Switch 2/64 and SAN Director 2/128 models), and use WWN zoning only.

If all the targets exist on the HP StorageWorks 2 GB switches (including Core Switch 2/64 and SAN Director 2/128 models), hardware enforced zoning will be active to protect them. This does not hardware-protect the initiators, but in general they do not offer resources anyway and will reject any probe attempt by other initiators.

Hardware-Enforced Zoning Configurations

Figure 2 shows a fabric with four non-overlapping zones, two defined by port number and two defined by WWN. These zones are hardware-enforced.

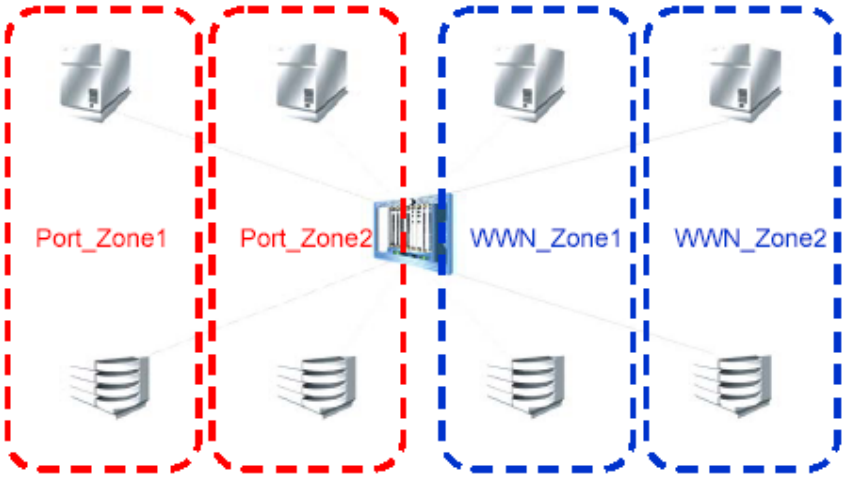


Figure 2: Nonoverlapping Hardware Enforced Zoning

Figure 3 shows a fabric with four zones, with overlap of like zone types.

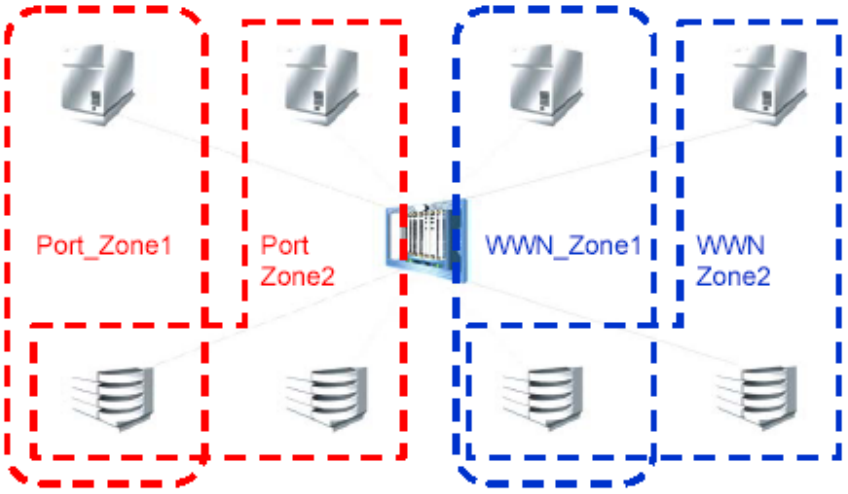


Figure 3: Overlapping Hardware-Enforced Zoning

Any zone comprising WWNs and *domain,port* entries on the 2 Gbit/sec platform relies on Name Server authentication as well as hardware-assisted (ASIC) authentication, which ensures that any PLOGI/ADISC/PDISC/ACC from an unauthorized device attempting to access a device it is not zoned with is rejected. 2 Gbit/sec switches always deploy the hardware assist, in any zone configuration. See Figure 4 and Figure 5.

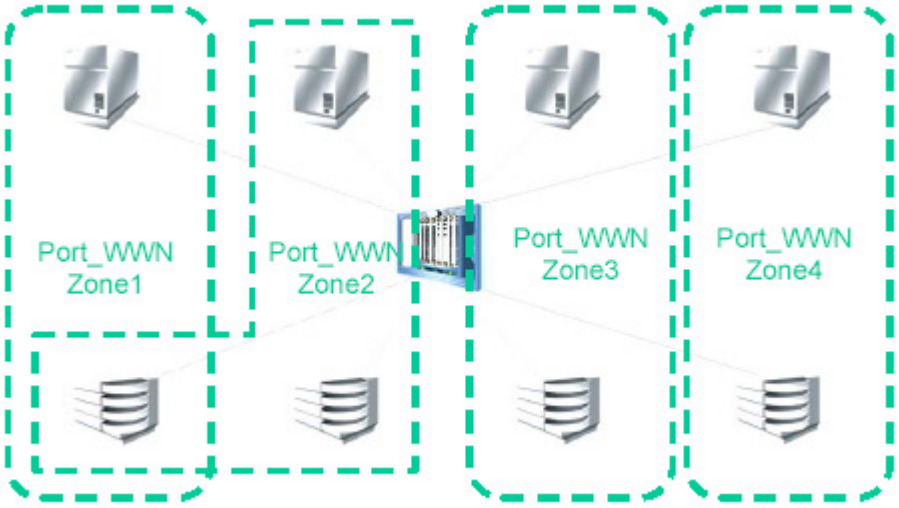


Figure 4: Zoning with Hardware Assist (Mixed Port and WWN Zones)

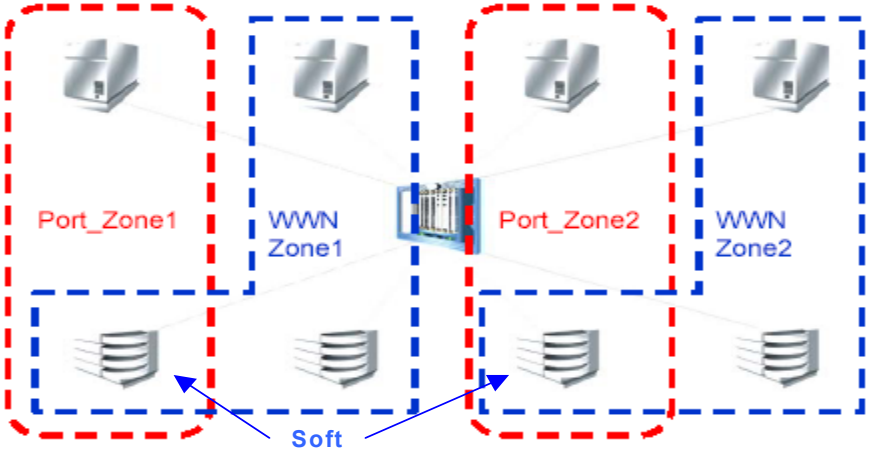


Figure 5: Overlapping Hardware Enforced Zoning with Soft Porting

In Figure 5, only the ports that are overlapped are software-enforced with hardware assist.

Rules for Configuring Zones

HP recommends using the following guidelines during zone configuration.

- Type of zoning (hard or soft)—If security is a priority, hard zoning is recommended.
- Use of aliases—The use of aliases is optional with zoning, and using aliases requires structure when defining zones. However, aliases aid administrators of a zoned fabric to understand the structure and context.
- Security requirements—Evaluate the security requirements of the fabric. If additional security is required, add Secure Fabric OS into the fabric.
- Interoperability fabric—If the fabric includes a Core Switch 2/64, SAN Director 2/128, or SAN Switch 2/32 and the user is supporting a third-party switch product, they are only able to use WWN zoning; other types of zoning, including QuickLoop, are not supported.
- QuickLoop—Evaluate if the fabric will have QuickLoop Fabric Assist (QLFA) or QuickLoop (QL) in it. If you are running Fabric OS v4.2, consider the following before creating and setting up QLFA zones:
 - QuickLoop Zoning, QuickLoop/QuickLoop zones cannot run on switches running Fabric OS v4.2.x. However, Fabric OS v4.2.x can still manage (create, remove, update) QuickLoop zones on any non-v4.2.x switch.
 - QuickLoop Fabric Assist, Fabric OS v4.2.x cannot have a Fabric Assist host directly connected to it. However, Fabric OS v4.x can still be part of a Fabric Assist zone if a Fabric Assist host is connected to a non-v4.2.x switch.
- Zone changes—Zone changes in a production fabric can result in a disruption of I/O under conditions in which an RSCN is issued as a result of a zone change and the HBA is unable to process the RSCN fast enough. Though RSCNs are a normal part of a functioning SAN, the pause in I/O might not be acceptable. For these reasons, you should perform zone changes only when the resulting behavior is predictable and acceptable. Changing HBA drivers can rectify the situation.
- Final verification—After changing or enabling a zone configuration, confirm that the nodes and storage are able to identify and access one another. Depending on the platform, you might need to reboot one or more nodes in the fabric with the new changes.

Fabric-Based Zoning

Zoning can be enabled in most storage units, in the switch fabric, and on the host. Fabric switches implement name server-based zoning, in which the zone members are identified by WWN or port location in the fabric.

Note: Refer to the *HP StorageWorks SAN Design Reference Guide* at the following location for general SAN zoning rules and requirements:
http://h18000.www1.hp.com/products/storageworks/san/documentation.html

Refer to the *HP StorageWorks Enterprise Backup Solutions Design Guide* at the following location for Enterprise Backup Solution (EBS) zoning recommendations:
http://h10025.www1.hp.com/ewfrf/wc/manualCategory?lc=en&cc=us&product=406722

- Storage-based zoning—Storage units typically implement LUN-based zoning, commonly referred to as *LUN masking*. Its basic function is to limit access to the LUNs on the storage port to the specific WWN of the server HBA.

 This form of zoning is needed in most SANs. It functions during the probe portion of the SCSI initialization. The server probes the storage port for a list of available LUNs and their properties. The storage system compares the WWN of the requesting HBA to the defined zone list and returns the LUNs assigned to the WWN. Any other LUNs on that storage port are not made available to the server.

- Fabric-based zoning—Fabric switches implement fabric-based zoning, in which the zone members are identified by WWN or port location in the fabric.

 Fabric-based Zoning is commonly referred to as *name server-based* or *soft* zoning (refer to "Software-Enforced Zoning" on page 33). With HP switches there might also be additional hardware enforcement of the zone (refer to "Hardware-Enforced Zoning" on page 33).

 When a device makes a query to the fabric name server, the name server determines to which zones the device belongs. It then returns to the requesting device information on all members of the zones present in the fabric. Devices in the zone are identified by node WWN, port WWN, or domain/port of the switch to which the device is connected.

- Host-based zoning—Host-based zoning can implement WWN or LUN masking.

 There are several approaches for implementing host-based zoning; all will work, in most cases. However, there are pros and cons to each form. The primary forms are: single HBA, grouping by application, grouping by operating system, port allocation, and no fabric zoning.

Zoning by Single HBA

Zoning by single host bus adapter (HBA) most closely re-creates the original SCSI bus. Each zone created has only one HBA (initiator) in the zone; each of the target devices is added to the zone. Typically, a zone is created for the HBA and the disk storage ports added. If the HBA also accesses tape devices, a second zone is created, with the HBA and associated tape devices in it.

In the case of clustered systems, it might be appropriate to have an HBA from each of the cluster members included in the zone. This is equivalent to having a shared SCSI bus between the cluster members and presumes that the clustering software can manage access to the shared devices.

In a large fabric, zoning by single HBA requires the creation of possibly hundreds of zones; however, each zone contains only a few members. Zone changes affect the smallest possible number of devices, minimizing the impact of an incorrect zone change.

Zoning by Application

Zoning by application typically requires zoning multiple, perhaps incompatible, operating systems into the same zones. This method of zoning creates the possibility that a minor server in the application suite could disrupt a major server, such as a Web server disrupting a data warehouse server.

Zoning by this method can also result in a zone with a large number of members, providing greater susceptibility to administrative errors, such as registered state change notifications (RSCNs) going out to a larger group than necessary.

Zoning by Operating System

Zoning by operating system allows multiple HBAs of the same operating system type to be grouped with the storage ports being accessed. This grouping helps keep the number of zones to a limited number in a given fabric. In some cases, the zone sizes may become too large; possibly, they can be broken into multiple zones within the operating system type, depending on the application. Zoning by operating system type can limit any disruptions or fabric change notifications (such as RSCNs) to a select few while at the same time retaining the advantage of grouping the HBAs of the same operating system type.

Also, zoning by operating system prevents HBAs of different operating systems from interacting with each other, potentially eliminating any incompatibilities between them.

Zoning by operating system is the preferred method. However, certain situations, such as when configuring for common server access to multiple storage types, may require zoning by HBA. For additional zoning requirement information, refer to the "Best Practices" chapter in the *HP StorageWorks SAN Design Reference Guide* at the following location:
http://h18000.www1.hp.com/products/storageworks/san/documentation.html

Zoning by Port Allocation

Zoning by port allocation should be avoided unless the administration team has very rigidly enforced processes for port and device allocation in the fabric. It does, however, provide some positive features. One feature for instance, when a storage port, server HBA, or tape drive is replaced, the change of WWN for the new device is of no consequence. As long as the new device is connected to the original port, it continues to have the same access rights. The ports on the edge switches can be pre-associated to storage ports, and control of the fan-in ratio can be established. With this pre-assigning technique, the administrative team cannot overload any one storage port by associating too many servers with it.

No Fabric Zoning

Using no fabric zoning is the least desirable zoning option because it allows devices to have unrestricted access on the fabric. Additionally, any device attached to the fabric, intentionally or maliciously, likewise has unrestricted access to the fabric.

This form of zoning should only be utilized in a small and tightly controlled environment, such as when host-based zoning or LUN masking is deployed.

The Effect of Zoning on RSCN Delivery

Whenever a change occurs in the name server, such as adding or removing a device from the fabric, a state change notification (RSCN) must be generated. In the absence of zoning, an RSCN is sent to all devices on the fabric, resulting in a brief but significant increase in unnecessary traffic.

After zoning is implemented, only the affected devices receive the RSCN. Table 3 lists possible state changes and their impact on various fabric operating systems.

Table 3: RSCN Delivery

State Change	FOS v2.x up to v2.5	FOS v3.x	FOS v2.6/3.0.2/4.x	FOS 3.1/4.1.x
Fabric reconfigure with no domain change	Fabric RSCN to all devices	Fabric RSCN to all devices	No RSCN	No RSCN
Switch disable	Fabric RSCN to all devices	Fabric RSCN to all devices	Domain RSCN to zoned devices only	No RSCN
Switch enable	Fabric RSCN to all devices	Fabric RSCN to all devices	No RSCN	No RSCN
Device power up	Device RSCN to zoned devices only	Aggregated[1] RSCN to zoned devices only	Aggregated RSCN to zoned devices only	Aggregated RSCN to affected zoned devices only
Device shutdown	Device RSCN to zoned devices only	Aggregated RSCN to zoned devices only	Aggregated RSCN to zoned devices only	Aggregated RSCN to affected zoned devices only
Zone change	Fabric RSCN to all devices	Fabric RSCN to all devices	Fabric RSCN to all devices	Fabric RSCN to affected devices only
Switchname/IP Address change	Domain RSCN to all devices	Domain RSCN to all devices	Domain RSCN to all devices	Domain RSCN to all devices

1. *Aggregated RSCN* combines multiple device RSCNs into one.

Using Zoning to Administer Security

Zones can provide controlled access to fabric segments and establish barriers between operating environments. They isolate systems with different uses, protecting individual systems in a heterogeneous environment; for example, when zoning is in secure mode, no merge operations occur.

Zoning is done on the primary Fabric Configuration Server (FCS). The primary FCS switch makes zoning changes and other security-related changes. The primary FCS switch also distributes zoning to all other switches in the secure fabric. All existing interfaces can be used to administer zoning (depending on the policies—refer to the *HP StorageWorks Secure Fabric OS 4.2.x User Guide* for information about security policies).

You must perform zone management operations from the primary FCS switch using a zone management interface, such as telnet or Advanced Web Tools. You can alter a zoning database, provided you are connected to the primary FCS switch.

When two secure fabrics join, the traditional zoning merge does not occur. Instead, a zoning database is downloaded from the primary FCS switch of the merged secure fabric. When E_Ports are active between two switches, the name of the FCS server and a zoning policy set version identifier are exchanged between the switches. If the views of the two secure fabrics are the same, the fabric's primary FCS server downloads the zoning database and security policy sets to each switch in the fabric. If there is a view conflict, the E_Ports are segmented due to incompatible security data.

As part of zoning architecture, you need to determine which of the two basic zoning architectures (hard or soft zoning) work best for your fabric. With time and planning, the basic hard zone configuration works for most sites. If a site has additional security needs, you need to add the additional layer of Secure Fabric OS, in addition to the standard zoning architecture.

Note: Secure Fabric OS requires the activation of an HP Security license and an HP Zoning license.

Zone Management

You can manage zoning using the following methods:

- Telnet, CLI
- Advanced Web Tools
- Fabric Manager
- Secure Shell
- Fabric Access (API) via third-party software

Refer to the *HP StorageWorks Fabric OS 4.2.x Procedures Guide* for information about managing zoning through telnet commands. Refer to the *HP StorageWorks Advanced Web Tools 4.2.x User Guide* for information about managing zoning through Advanced Web Tools. Refer to the *HP StorageWorks Fabric Manager 4.1.1 User Guide* for information about managing zoning through Fabric Manager.

Zoning commands are executed under the transaction model. A working copy of the defined configuration is created at the start of a transaction. Each zoning command is executed from the working copy.

When a transaction is opened, all new zoning information is placed in a transactional buffer. The new changes are not applied to the fabric until the transaction is closed. A transaction is aborted when another switch closes its transaction or by issuing the command `cfgTransAbort`. When a transaction is closed, all new and existing zoning information is applied to the fabric and saved to flash memory.

Zone configuration is managed on a fabric basis. Zoning can be implemented and administered from any switch in the fabric that has a Zoning license enabled. When a change in the configuration is saved, enabled, or disabled per the transactional model, it is automatically (by closing the transaction) distributed to all switches in the fabric, preventing a single point of failure for zone information.

Table 4 lists the telnet commands used in zoning management.

Table 4: Commands Used for Managing Zoning

Command	Description
Commands to open a transaction	
aliCreate	Create a zone alias.
aliDelete	Delete a zone alias.
aliRemove	Remove a member from a zone alias.
cfgAdd	Add a member to a configuration.
cfgClear	Clear all zone configurations.
cfgCreate	Create a zone configuration.
cfgDelete	Delete a zone configuration.
cfgRemove	Remove a member from a configuration.
fazoneAdd	Add a member to a fabric assist zone.
fazoneCreate	Create a fabric assist zone.
fazoneDelete	Delete a fabric assist zone.
fazoneRemove	Remove a member from a fabric assist zone.
zoneAdd	Add a member to a zone.
zoneCreate	Create a zone.
zoneDelete	Delete a zone.
zoneRemove	Remove a member from a zone.
Commands to close a transaction	
cfgDisable	Disable a zone configuration and save zone configuration in flash.
cfgEnable	Enable a zone configuration.
cfgSave	Save zone configurations in flash.
Informative commands (do not impact a transaction)	
aliShow	Print aliases.
cfgShow	Print configuration information.
zoneShow	Print zone configuration information.

Advanced Performance Monitoring

3

This chapter contains the following information:

- Overview of Advanced Performance Monitoring, page 48
- Types of Performance Monitoring, page 49
- Operating Advanced Performance Monitoring with Advanced Web Tools, page 54
- Performance Monitoring with Telnet Commands, page 58

Overview of Advanced Performance Monitoring

Advanced Performance Monitoring is an optionally licensed product used for monitoring the performance of networked storage resources. This tool helps reduce overprovisioning while enabling SAN performance tuning and increasing administrator productivity.

Advanced Performance Monitoring provides SAN performance monitoring through an end-to-end monitoring system providing:

- Increased end-to-end visibility into the fabric.
- More accurate reporting for service-level agreements and charged access applications.
- Shortened troubleshooting time.
- Better capacity planning.
- Increased productivity via preformatted and customized screens and reports.

Advanced Performance Monitoring is administered through either telnet commands or Advanced Web Tools.

The Advanced Performance Monitoring product:

- Measures the bandwidth consumed by individual routes (host-target pairs).
- Provides device performance measurements by port, AL_PA, and LUN.
- Reports CRC error measurement statistics.
- Compares IP versus SCSI traffic on each port.
- Includes a wide range of predefined reports.
- Allows you to create customized user-defined reports.

Types of Performance Monitoring

Using Advanced Performance Monitoring, you can track the following:

- Number of CRC errors for AL_PA devices
- Number of words received and transmitted in Fibre Channel frames with a defined SID/DID pair
- Number of frames with CRC errors received at the port with a defined SID/DID pair
- Number of times a particular filter pattern in a frame is transmitted by a port

Telnet commands and Advanced Web Tools provide access to three types of performance monitoring:

- AL_PA monitoring
- End-to-end monitoring
- Filter-based monitoring

AL_PA Monitoring

AL_PA performance monitoring tracks and displays the number of CRC errors that have occurred on frames sent from each AL_PA on a specific port. AL_PA-based performance monitoring does not require explicit configuration. The switch hardware and firmware automatically monitors CRC errors for all valid AL_PAs.

End-to-End Monitoring

End-to-end monitoring provides information regarding performance between the source (SID) and destination (DID) on a fabric or a loop. Up to eight SID-DID pairs can be specified per port. For each of the SID-DID pairs, the following information is available:

- Number of fibre channel words received by the port for the SID-DID pair (RX_COUNT)
- Number of fibre channel words transmitted from the port for the SID-DID pair (TX_COUNT)
- Number of frames with CRC errors received at or transmitted from the port for the SID-DID pair (CRC_COUNT)

To enable end-to-end performance monitoring, you must configure an end-to-end monitor on a port, specifying the SID-DID pair. The monitor counts only those frames with matching SID and DID. Each SID or DID has three fields, listed in the following order:

- Domain ID (DD)
- Area ID (AA)
- AL_PA (PP)

For example, the SID 0x118a0f has domain ID *0x11*, area ID *0x8a*, and AL_PA *0x0f*. (The prefix "0x" denotes a hexadecimal number.)

The monitor counts the number of words received, number of words transmitted, and number of CRC errors detected in frames qualified using either of the following two conditions:

- For frames received at the port (with end-to-end monitor installed), the frame SID is the same as "SourceID" and the frame DID is the same as "DestID." Both RX_COUNT and CRC_COUNT are updated accordingly.
- For frames transmitted from the port (with end-to-end monitor installed), the frame DID is the same as "SourceID" and the frame SID is the same as "DestID." TX_COUNT and CRC_COUNT are updated accordingly.

Where to Add End-to-End Monitors

Depending on the application, any port along the routing path can be selected for such monitoring.

Figure 6 shows two devices:

- Host A, connected to domain 0x05, switch area port 12 (area ID 0x0c), AL_PA 0x00 on Switch X
- Dev B, connected to domain 0x11, switch area port 30 (area ID 0x1e), AL_PA 0xef on Switch Y

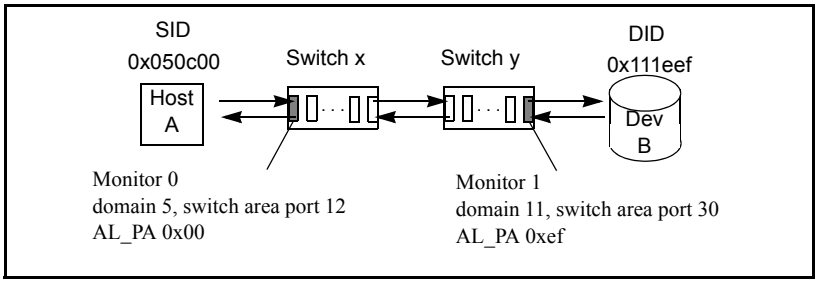


Figure 6: Setting End-to-End Monitors on a Port

To monitor the traffic from Host A to Dev B, add a monitor to slot 1, port 12, specifying 0x050c00 as the SID and 0x111eef as the DID. To monitor the traffic from Dev B to Host A, add a monitor to slot 2, port 14, specifying 0x111eef as the SID and 0x050c00 as the DID.

Monitor 0 counts the frames that have an SID of 0x050c00 and a DID of 0x111eef. For monitor 0:

RX_COUNT =the number of words from Host A to Dev B
TX_COUNT=the number of words from Dev B to Host A
CRC_COUNT=the number of frames in both directions with CRC errors

Monitor 1 counts the frames that have an SID of 0x111eef and a DID of 0x050c00. For monitor 1:

RX_COUNT=the number of words from Dev B to Host A
TX_COUNT=the number of words from Host A to Dev B
CRC_COUNT=the number of frames in both directions with CRC errors

Note: End-to-end performance monitoring monitors traffic on the receiving port respective to the SID only. In Figure 6, if you add a monitor to slot 2, port 14 specifying Dev B as the SID and Host A as the DID, no counters, except CRC, will be incremented.

Figure 7 shows several switches and the proper ports on which to add performance monitors for a specified SID-DID pair.

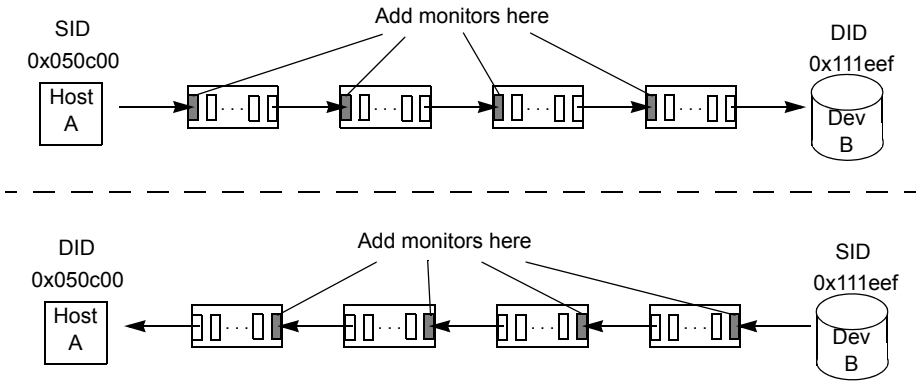


Figure 7: Proper Placement of End-to-End Performance Monitors

End-to-end monitors count the number of words in Fibre Channel frames that match a specific SID/DID pair. If you want to match only part of the SID or DID, you can set a mask on the port to compare only certain parts of the SID or DID. With no mask set, the frame must match the entire SID and DID to trigger the monitor. By setting a mask, you can choose to have the frame match only one or two of the three fields (domain ID, Area ID, and AL_PA) to trigger the monitor.

Note: You can set only one mask per port. When you set a mask, all existing end-to-end monitors are deleted.

For example, you can set a mask to compare the AL_PA fields on the SID and DID in all frames (transmitted and received) on a specific port; then, the frame SID and DID must match only the AL_PA portion of the specified SID-DID pair.

Each port can have only one EE mask. The mask is applied to all end-to-end monitors on the port. Individual masks for each monitor on the port cannot be specified. The default EE mask value upon power-on is a mask for everything; thus, the SID and the DID must match completely in order to increment the counters.

Filter-Based Monitoring

Filter-based monitoring counts the number of times a frame with a particular pattern is transmitted by a port. Filter-based monitoring is achieved by configuring a filter for a particular purpose. The filter can be a standard filter (for example, a read command filter that counts the number of read commands transmitted by the port) or a user-defined filter that you customize for your particular use. The maximum number of filters is eight per port; however, depending on your combination of standard and user-defined filters, that number might be reduced.

Standard Filter-Based Monitors

Using the standard filters, you can collect the following fibre channel frame statistics:

- Number of SCSI read, write, or read/write commands
- Number of SCSI traffic frames
- Number of IP traffic frames

User-Defined Filter-Based Monitors

In addition to the standard filters (read, write, read/write, frame count, and IP count), you can create custom filters to qualify frames to gather statistics to fit your needs.

To define a custom filter, use the `perfaddusermonitor` telnet command.

Operating Advanced Performance Monitoring with Advanced Web Tools

Use Advanced Web Tools to view, customize, or monitor performance by:

- Viewing predefined reports for AL_PA, end-to-end, and filter-based performance monitoring.
- Creating user-definable reports.
- Viewing performance canvas for application-level or fabric-level views.
- Accessing configuration editor (save, copy, edit, and remove multiple configurations).
- Saving persistent graphs across reboots (saves parameter data across reboots).

For detailed information, refer to the *HP StorageWorks Advanced Web Tools 4.2.x User Guide*.

Predefined Performance Graphs

Advanced Web Tools provides graphs to simplify performance monitoring. A wide range of end-to-end fabric, LUN, device, and port metrics graphs are included. Table 5 lists the performance graphs available. You can access the basic monitoring graphs on switches that do not have an Advanced Performance Monitoring license activated. The advanced monitoring graphs give more detailed performance information to help you manage your fabric.

Table 5: Predefined Performance Graphs

Basic Monitoring Graphs	Description
Port Throughput	Displays the performance of a port in bytes per second for frames received and transmitted.
Switch Aggregate Throughput	Displays the aggregate performance of all ports on a switch.
Blade Aggregate Throughput	Displays the aggregate performance of all ports on a blade.
Switch Throughput Utilization	Displays the port throughput at the time the sample is taken.
Port Error	Displays a line of CRC errors for a given port.
Switch Percent Utilization	Displays the percentage of usage of a chosen switch at the time the sample is taken.
Port Snapshot Error	Displays the CRC error count between sampling periods for all the ports on a switch.
Advanced Monitoring Graphs	**Description**
SID/DID Performance	Displays the traffic between SID and DID pair on the switch being managed.
SCSI vs. IP Traffic	Displays percentage of SCSI versus IP frame traffic on each individual port.
AL_PA Errors	Displays CRC errors for a given port and a given AL_PA.
SCSI Commands by port and LUN (R, W, R/W)	Displays the total number of read/write commands on a given port and read/write commands to a specific LUN.

Figure 8 shows how to access the list of Advanced Performance Monitoring graphs using Advanced Web Tools.

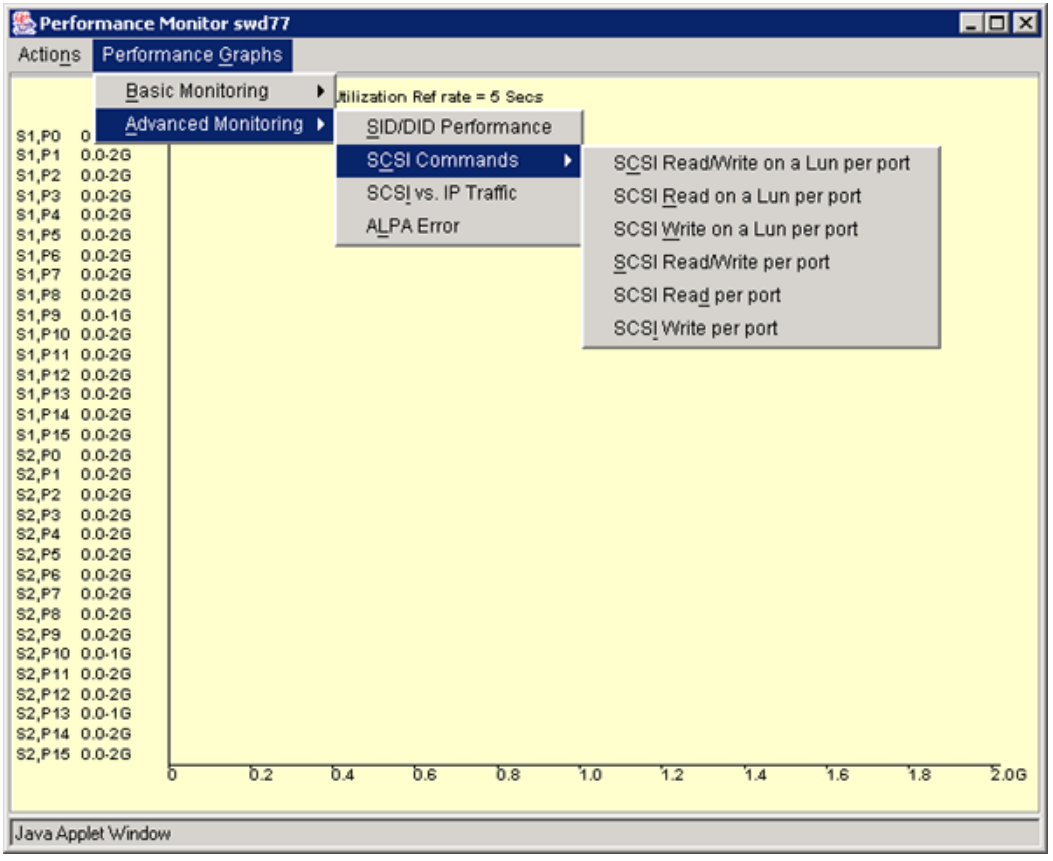


Figure 8: Accessing Performance Graphs

User-Defined Graphs

You can modify the predefined graphs based on parameter fields such as SID/DID, LUN, AL_PA, and port to create your own customized graphs. These user-defined graphs can be added and saved to canvas configurations.

Using Advanced Web Tools, Advanced Performance Monitoring allows you to set up a *canvas* of performance graphs. The canvas can hold up to eight graphs per window, as shown in Figure 9. Multiple canvases can be set up for different users or different scenarios.

Up to 20 individual canvases, each with up to eight graphs, can be saved. Each canvas is saved with a name and an optional brief description.

In addition to the graphs, the Performance Monitoring Resource Usage Display shows which filter slots have been used for each port, and which are available. A color-coded interface facilitates troubleshooting. See Figure 9.

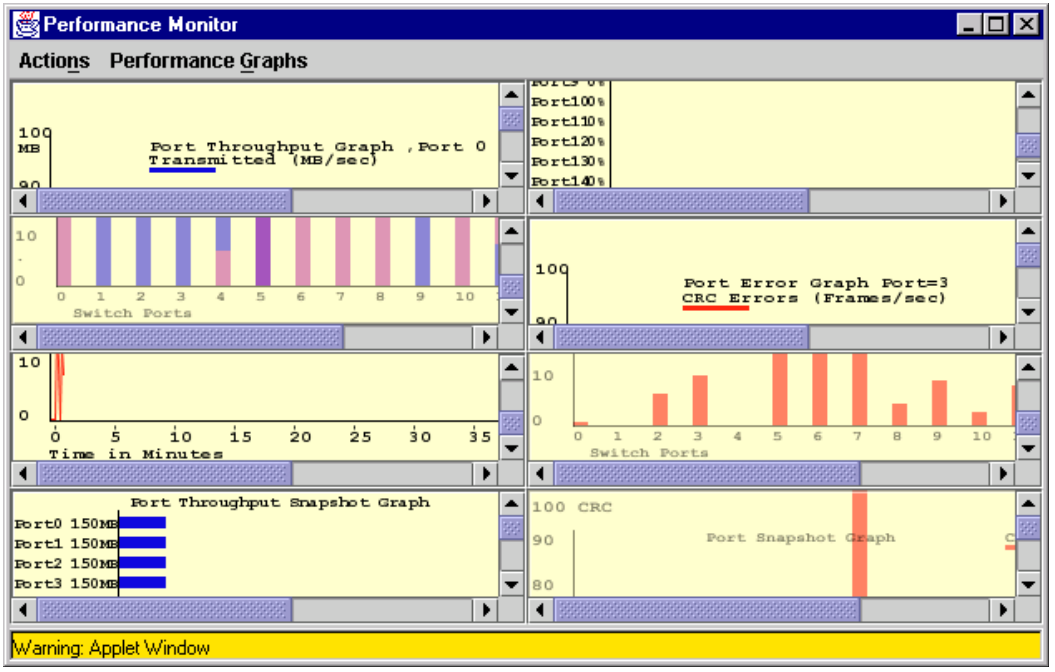


Figure 9: Eight Performance Monitoring Graphs

Performance Monitoring with Telnet Commands

Table 6 lists the telnet commands for Advanced Performance Monitoring. For additional information, refer to the *HP StorageWorks Fabric OS 4.2.x Reference Guide*.

Table 6: APM Telnet Commands

Command	Description
AL_PA Monitoring	
perfshowalpacrc	Displays the CRC error count for all AL_PA devices or a single AL_PA on a specific port. The port must be an active L_Port.
perfclralpacrc	Clears the CRC error count for a specific AL_PA or for all AL_PA devices on a port. **Note:** In Fabric OS versions 3.x and 4.x, issuing the `portStatsClear` command on a port also results in all AL_PA-based CRC error counters being cleared for all the ports in the same quad.
End-to-End Monitoring	
perfaddeemonitor	Adds an end-to-end monitor to a port.
perfsetporteemask	Sets a mask for all end-to-end monitors of a port.
perfshowporteemask	Displays the current end-to-end mask of a port.
perfshoweemonitor	Displays all the end-to-end monitors defined on a port.
perfdeleemonitor	Deletes an end-to-end monitor on a port.
perfcleareemonitor	Clears statistics counters for all or a specified end-to-end monitor on a port.
Filter-Based Monitoring	
perfaddreadmonitor	Adds a filter-based monitor to count the number of SCSI read commands.
perfaddwritemonitor	Adds a filter-based monitor to count the number of SCSI write commands.
perfaddrwmonitor	Adds a filter-based monitor to count the number of SCSI read and write commands.
perfaddscsimonitor	Adds a filter-based monitor to count the number of SCSI traffic frames.
perfaddipmonitor	Adds a filter-based monitor to count the number of IP traffic frames.
perfshowfiltermonitor	Displays all the filter-based monitors configured on a port.
perfaddusermonitor	Defines a custom filter.
perfdelfiltermonitor	Deletes a specific monitor.

Table 6: APM Telnet Commands (Continued)

Command	Description
perfclearfiltermonitor	Clears statistics counters for all or a specified filter-based monitor.
Other Commands	
perfcfgsave	Saves the current end-to-end and filter monitor configuration settings into flash memory.
perfcfgrestore	Restores the saved monitor configuration from flash memory.
perfcfgclear	Clears the previously saved performance monitoring configuration settings from flash memory.

ISL Trunking

4

This chapter provides the following information:

- Overview of ISL Trunking, page 62
- ISL Trunking Architecture, page 63
- Setting Up ISL Trunking in a Fabric, page 71
- Managing ISL Trunking, page 74
- Managing ISL Trunking, page 74
- Trunk Information Tab in the Switch Admin Window of Advanced Web Tools, page 75
- Frequently Asked Questions About ISL Trunking, page 77

Overview of ISL Trunking

ISL Trunking is an optionally licensed product available on all hp StorageWorks 2 Gb switches.

ISL Trunking is licensed on a per-switch basis. It optimizes network performance by forming trunking groups that can distribute traffic across the shared bandwidth of all the interswitch links (ISLs) in the trunking group. It is compatible with both short wavelength (SWL) and long wavelength (LWL) fiber optic cables and transceivers.

ISL Trunking allows up to four ISL connections between two switches to merge logically into a single link. This enables traffic to be routed through any available ISL in the group rather than being restricted to a specific, potentially congested ISL. ISL Trunking distributes traffic dynamically across the merged ISLs at the fibre channel frame level while preserving in-order delivery of the frames.

ISL Trunking Architecture

The ISL Trunking software identifies and constructs trunking groups as soon as the ISL Trunking license is activated. The ISLs and ports that participate in trunking groups are referred to as *trunking ISLs* and *trunking ports*.

ISL Trunking makes it possible to accomplish the same fabric performance with fewer ISLs, resulting in simplified fabric design and management, lowered cost of ownership, increased fabric performance, and increased data availability.

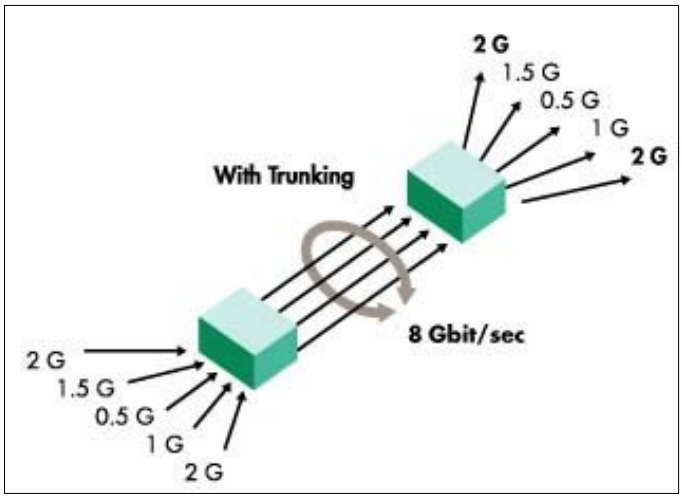


Figure 10: Distribution of Traffic Over an ISL Trunking Group

Figure 10 illustrates how trunking can result in more throughput by avoiding congestion. In this example, the data available for transmission is distributed over the four ISLs with no congestion, since it is below the total 8 Gbit/sec capacity of the combined ISLs. In a fabric that does not have trunking capability, some paths would be congested and other paths underutilized.

The following sections describe the traffic routing protocol, trunking groups, trunking ISLs, and trunking ports.

Traffic Routing

The same routing protocol, Fabric Shortest Path First (FSPF), is used both with and without ISL Trunking. FSPF directs traffic along the shortest path between source and destination, based on the link cost, and makes it possible to detect link failures, determine the shortest route for traffic, update the routing table, provide fixed routing paths within a fabric, and maintain correct frame order.

ISL Trunking reduces or eliminates situations that require static traffic routes and individual ISL management in order to achieve optimal performance. If any static routes are specified using the `urouteconfig` command, the command overrides the dynamic load sharing employed by ISL Trunking, and performance might decrease accordingly.

ISL Trunking uses *deskew values*, which are a measure of the latency differences between ISL cables in the same group, to provide in-order traffic delivery. Deskew values are automatically calculated at the creation of each new trunking group.

Trunking Groups

A trunking group can contain two to four ISLs. If any eligible ISLs exist, trunking groups automatically form when the ISL Trunking license is activated. Multiple trunking groups can exist between the same two switches, and each switch can support as many trunking groups as allowed by the available ports.

Trunking ISLs

The first ISL discovered in any trunking group is automatically designated as the trunking *master ISL*. Each trunking group has one master ISL and up to three slave ISLs. The master ISL cannot be specified by the user.

If a master ISL fails or is removed, there might be a pause in the I/O associated with that trunk as a new master ISL is designated and traffic is redistributed. If data frames were in flight during failure of the master ISL, some frames might be lost, depending on the Fibre Channel class: Class 2 frames are resent and Class 3 frames are dropped. The impact of any dropped frames on the I/O depends on the host, application, storage, and host bus adapter (HBA) in use.

If a slave ISL fails or is removed, the traffic is redistributed without interruption over the remaining members in the trunking group.

The difference in cable length between the shortest ISL and the longest ISL in a trunking group must be less than 400 meters; less than 30 meters is recommended.

Trunking Ports

The ports on either end of the master ISL are called trunking *master ports*. These ports work together to direct traffic over the trunked ISLs. All other ports are called trunking *slave ports*.

All trunking ports must meet the following criteria:

- All ports in the same trunking group must reside in the same quad; *quads* are groups of four adjacent ports marked at intervals on each switch. The first four ports on a switch compose the first quad, and every group of four adjacent ports thereafter compose another quad. Figure 11 through Figure 17 show the quads for various switch models.
- The port speed must be set to autonegotiate or to 2 Gbit/sec (the default speed is autonegotiate).
- If Extended Fabrics is in use, the ports must be in L0 mode. Trunking is supported for normal E_Ports (referred to as L0 in the `portcfglongdistance` command) with LWL media up to 5km at the full speed permitted by the link. With LWL media, the throughput begins to fall off beyond 5km, due to normal latency effects. ISL Trunking does not support the LE, L1, or L2 `portcfglongdistance` modes. For information about these modes and Extended Fabrics in general, refer to the *HP StorageWorks Fabric OS 4.2.x Procedure Guide*.

Potential trunking groups are marked on the front of the switch, according to quad location.

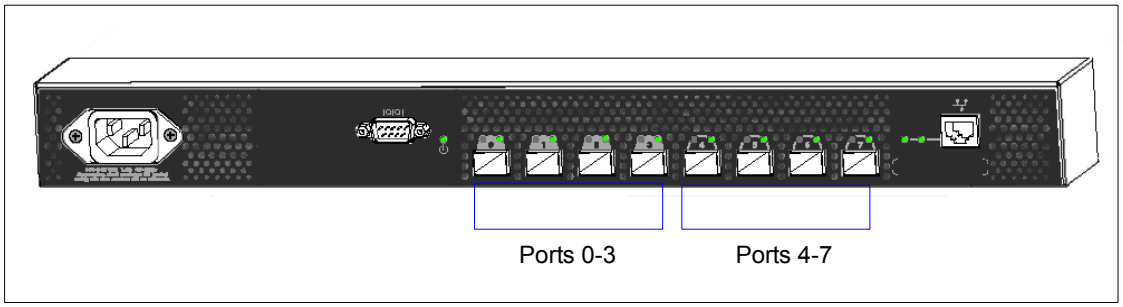


Figure 11: Port Groupings on a SAN Switch 2/8 EL

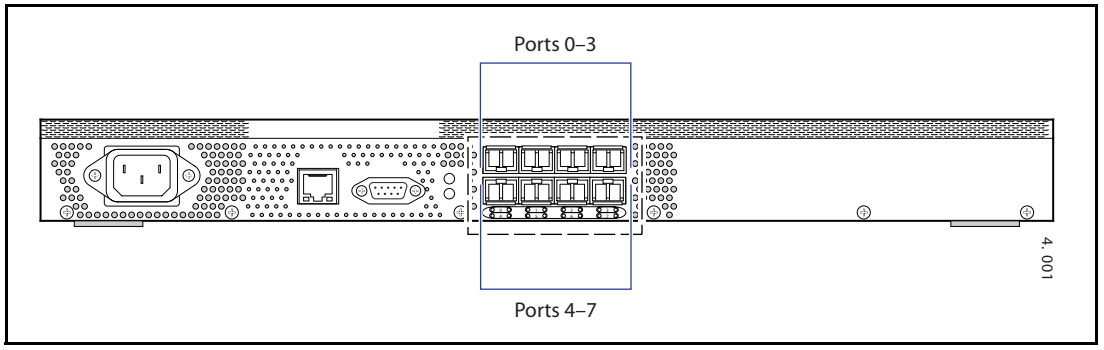


Figure 12: Quad and Port Groupings on a SAN Switch 2/8V

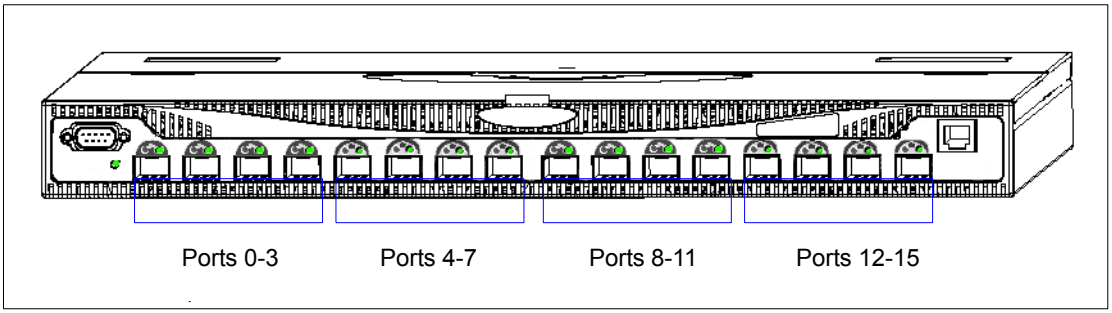


Figure 13: Quad and Port Groupings on a SAN Switch 2/16

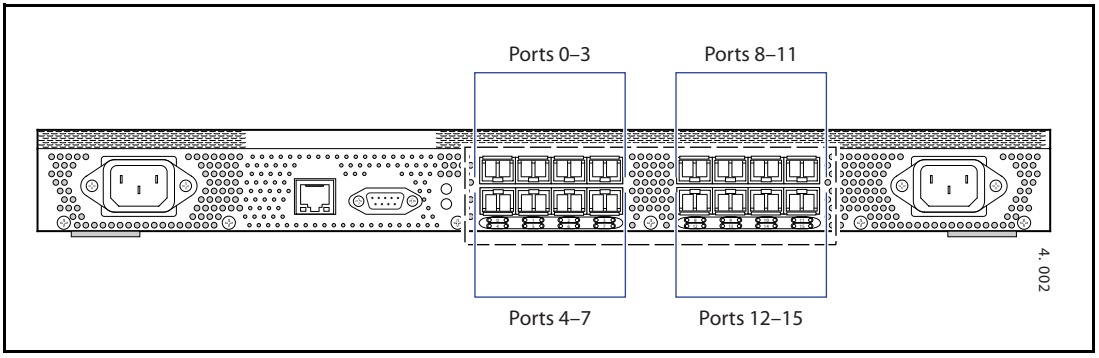


Figure 14: Quad and Port Groupings on a SAN Switch 2/16V

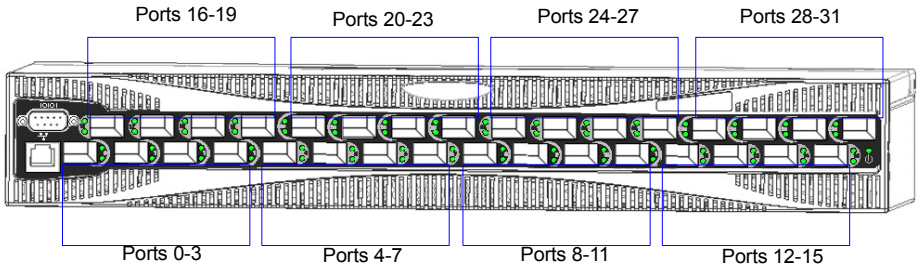


Figure 15: Quad and Port Groupings on a SAN Switch 2/32

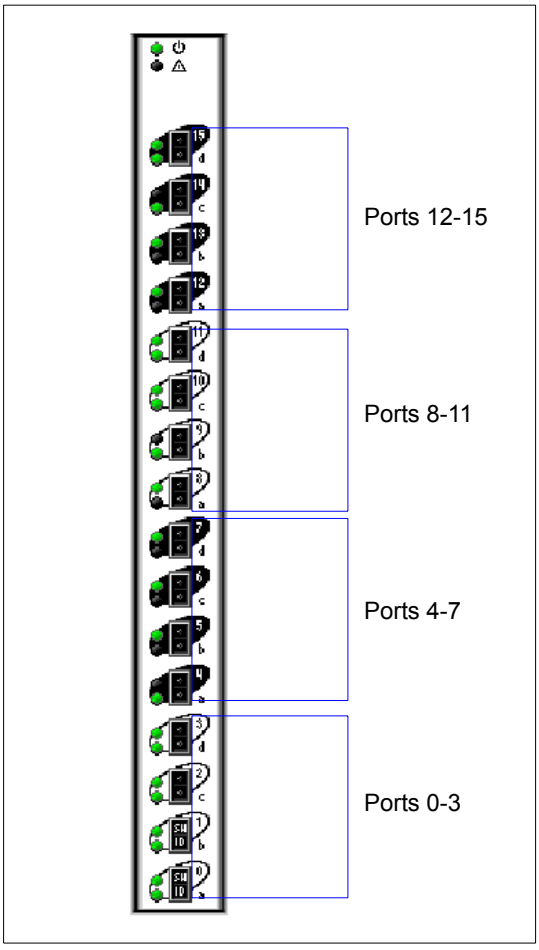


Figure 16: Port Groupings on the Port Card on a Core Switch 2/64

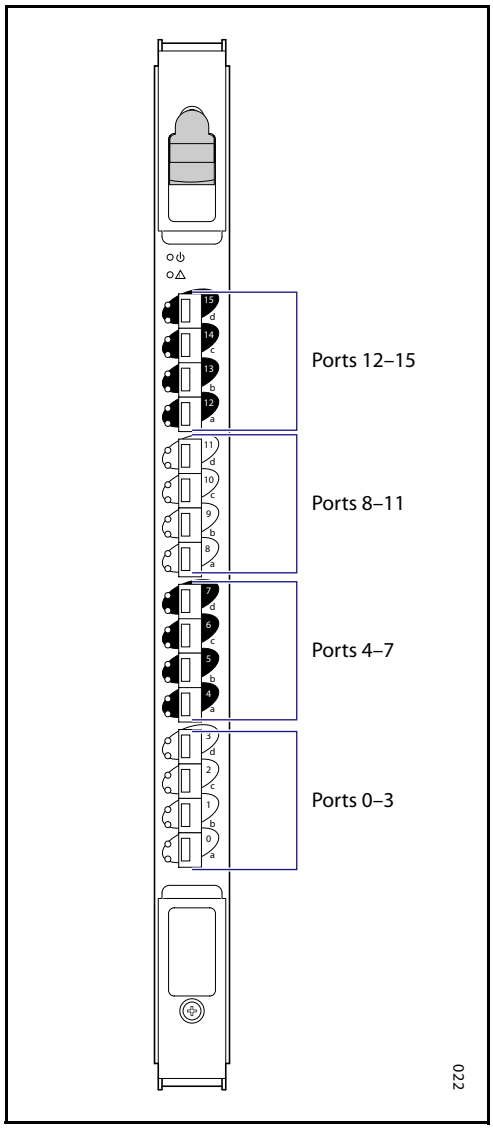


Figure 17: Port Groupings on the Port Card on a SAN Director 2/128

Setting Up ISL Trunking in a Fabric

The ISL Trunking feature is provided with the Fabric OS and can be activated by entering a license key, available from the switch supplier. When the ISL Trunking license is activated, trunking is automatically implemented for any eligible ISLs.

A license must be activated on each switch that will participate in trunking. For the Core Switch 2/64, a single license key enables the feature on both logical switches.

To use ISL Trunking in the fabric, the fabric must be designed to allow trunking groups to form. To identify the most useful trunking groups, evaluate the traffic patterns before designing/redesigning the fabric.

Note: ISL Trunking is enabled by default for each port on the switch as shipped. For instructions on disabling and enabling trunking capability for individual ports or all the ports on the switch, refer to the *HP StorageWorks Fabric OS 4.2.x Procedures User Guide*.

Evaluating Data Traffic Patterns

Traffic patterns can be monitored using the `portperfshow` command, the Advanced Performance Monitoring feature, the Fabric Watch feature, or a combination:

- Use the `portperfshow` command to record the traffic volume for each port over time to identify the congested paths that would benefit from the implementation of trunking groups. You can also use this command to identify frequently dropped links, so that troubleshooting can be performed and the links to trunking groups can be reinstated as necessary.
- Use Advanced Performance Monitoring to monitor traffic flow and to view the impact of different fabric configurations on performance.

 For instructions on using Advanced Performance Monitoring, refer to the *HP StorageWorks Fabric OS 4.2.x Procedures User Guide*.
- Use Fabric Watch to monitor traffic flow through specified ports on the switch and send alerts when the traffic exceeds or drops below configurable thresholds. This allows the administrator to monitor changes in traffic patterns and adjust the fabric design accordingly by adding, removing, or reconfiguring ISLs and trunking groups.

For instructions on configuring Fabric Watch thresholds and alerts, refer to the *HP StorageWorks Fabric Watch 4.2.x User Guide*.

Designing the Fabric to Optimize Use of ISL Trunking

ISL Trunking can be used to simplify SAN design and improve performance. When designing the SAN, consider the following recommendations (in addition to the standard guidelines for SAN design):

- Evaluate the traffic patterns within the fabric.

 This allows your fabric to implement trunking groups that will help to optimize its performance. Refer to "Evaluating Data Traffic Patterns" on page 71 for more information.

- Place trunking-capable switches adjacent to each other.

 This maximizes the number of trunking groups that can form. If you are using a core/edge topology, place trunking-capable switches at the core of the fabric and any switches that are not trunking-capable at the edge of the fabric.

- Activate an ISL Trunking license on each switch that is expected to participate in a trunking group.

- Verify that the lengths of the ISLs in the group do not differ by more than the recommended values.

 The lengths of the ISLs in the group should differ by less than 30 meters (recommended), and *must* differ by less than 400 meters.

 This is because large differences in latency decrease the efficiency of load-sharing. If cable lengths differ by 400 meters or more, the trunking group forms only for the ISLs that have lengths that differ by less than 400 meters.

- When connecting two switches with two or more ISLs, ensure that all trunking requirements are met to allow a trunking group to form.

- Determine the optimal number of trunking groups between each set of linked switches, depending on traffic patterns and port availability.

 The goal is to avoid traffic congestion without unnecessarily using ports that could be used to attach other switches or devices. Consider these points:

 — Each physical ISL uses two ports that could otherwise be used to attach node devices or other switches.

 — Trunking groups can be used to resolve ISL oversubscription if the total capability of the trunking group is not exceeded.

- Consider how the addition of a new path will affect existing traffic patterns:
 - A trunking group has the same link cost as the master ISL of the group, regardless of the number of ISLs in the group. This allows slave ISLs to be added or removed without causing data to be rerouted, because the link cost remains constant.
 - The addition of a path that is shorter than existing paths causes traffic to be rerouted through that path.
 - The addition of a path that is longer than existing paths might not be useful because the traffic will choose the shorter paths first.
- Plan for future bandwidth addition to accommodate increase in traffic.

 For trunking groups over which traffic is likely to increase as business requirements grow, consider leaving one or two ports in the quad available for future nondisruptive addition of bandwidth.
- Consider creating redundant trunking groups where additional ports are available or paths are particularly critical.

 This helps to protect against oversubscription of trunking groups, multiple ISL failures in the same group, and the rare occurrence of an ASIC failure.
- To provide the highest level of reliability, deploy trunking groups in redundant fabrics to further ensure ISL failures do not disrupt business operations.

Note: SWL and LWL fiber optic cables and transceivers can be used in the same trunking group.

Managing ISL Trunking

ISL Trunking provides a number of management options, including the following:

- Enabling and disabling trunking on entire switches or individual ports.
- Setting port speeds for entire switches or individual ports.

 If trunking is enabled, the only supported speeds are 2 Gbit/sec and autonegotiate. If trunking is not enabled, 1 Gbit/sec is also supported.
- Displaying trunking information.

 Advanced Web Tools lists each trunking group on the switch, with the master and slave ports for each group, as shown in Figure 18.
- Debugging any trunking failures.

You can manage ISL Trunking using a telnet or serial command line interface, Advanced Web Tools, or Fabric Manager. You can use standard SNMP applications to view but not to write trunking information.

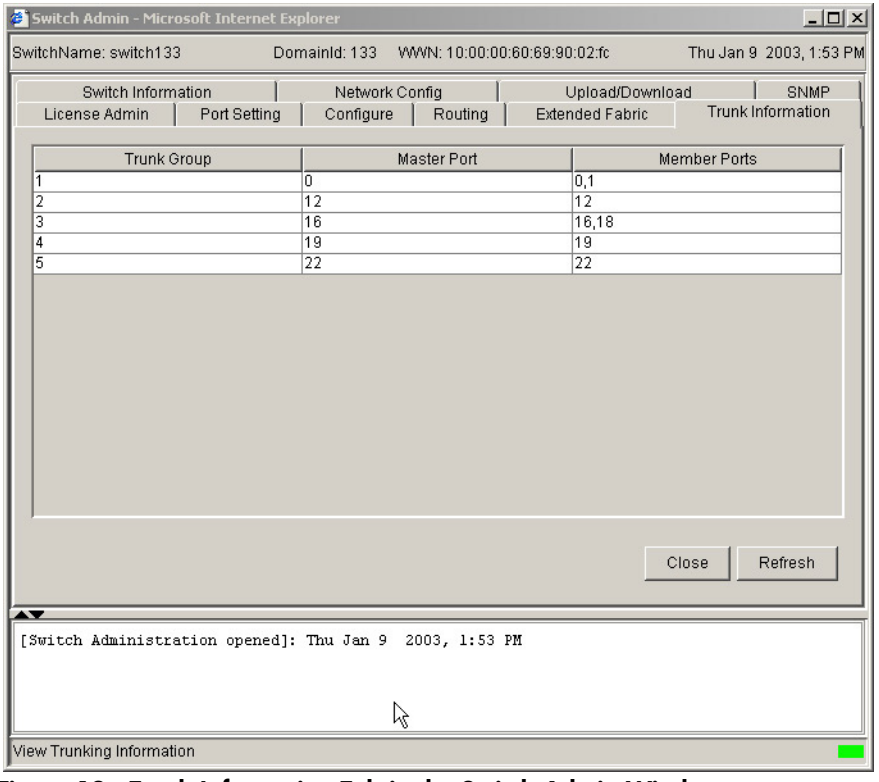


Figure 18: Trunk Information Tab in the Switch Admin Window of Advanced Web Tools

Commands Related to ISL Trunking

ISL Trunking includes several commands that can be used to configure settings related to trunking, such as port speeds and the ability to enable and disable trunking. Table 7 provides summary information about the commands that relate to ISL Trunking.

Table 7: Fabric OS Commands Related to ISL Trunking

Command	Function	Instructions
`portcfgspeed`	Use to specify the port speed.	Enter on the switch with the port to be configured.
`portcfgtrunkport`	Use to enable or disable trunking for a port.	Enter on the switch with the port to be configured.
`switchcfgspeed`	Use to set all ports of the switch to a particular speed.	Enter on the switch to specify port speeds.
`switchcfgtrunk`	Use to enable or disable trunking for all the ports of a switch.	Enter on the switch to enable or disable ISL Trunking.
`trunkdebug`	Use to debug a trunk link failure.	Enter on the switch to debug a trunk failure.
`trunkshow`	Use to display trunking information.	Enter on the switch to show ISL Trunking information.

Frequently Asked Questions About ISL Trunking

Table 8 provides answers to some frequently asked questions regarding ISL Trunking.

Table 8: Frequently Asked Questions about ISL Trunking

Frequently Asked Question	Answer
Does ISL Trunking replace Dense Wavelength Digital Multiplexing (DWDM)?	No, DWDM is a ring topology with a different function than trunking. If a DWDM ISL fails, the traffic is rerouted over alternate routes, changing the data path.
Is a trunking master ISL the same as the principal ISL?	No, the roles are different, although they might happen to apply to the same ISL. "Trunking master ISL" applies to the role of directing traffic over a trunking group. "Principal ISL" applies to an ISL that is used to communicate with the principal switch. This ISL is the communication path from one switch to the principal switch among many other ISLs.
Is it possible to create a trunk between a switch and an SAN device, such as host or storage?	No; ISL Trunking is only supported for interswitch links.
Is there a limit on the number of trunking groups on one switch?	No, the number of trunking groups that can be implemented on a switch is limited only by the number of available ports.
Are trunks automatically established when the ISL Trunking license is activated?	Yes, if eligible ISLs exist. Trunking capability is enabled by default on each port.
What happens if a slave ISL fails?	The traffic is redistributed over the remaining ISLs in the group.
What happens if a master ISL fails?	A new master ISL is designated and traffic is redistributed. If any in-transit frames are lost, there is a brief pause in the I/O.
Should port statistics be the same across all participating ISLs within a trunk?	No; although port statistics are usually fairly evenly balanced, they can vary with payload variations at the frame level.
Which Extended Fabric modes are supported?	"L0" mode, which is the default mode. If the ports in the potential trunking group use any other modes, the trunking group does not form.
Is trunking supported for 1 Gbit/sec?	No, trunking requires 2 Gbit/sec capacity.

Extended Fabrics

5

Extended Fabrics is an optionally licensed product that enables interswitch links (ISLs) to extend up to 100 km. This is achieved by optimizing the internal buffering algorithm used by HP StorageWorks switches. It provides maximum buffering between E_Ports connected over an extended distance. The buffer reconfiguration results in line speed performance of 2 Gbit/sec for switches interconnected at up to 60 km, and of 1 Gbit/sec for switches interconnected at up to 100km.

The Extended Fabric feature achieves long-distance connections by allocating more frame buffers for fibre channel traffic. Long-distance connections require more frame buffers than regular ISL connections. The greater the distance, the more frame buffers required. If the long-distance port is part of a quad, this limits the buffer space left over for the remaining ports in the quad, which must, therefore, be configured appropriately.

When configuring long distance ISLs, make sure to balance the need between long distance ISL connections and core-to-edge ISL connections within a switch. Configuring long distance ISLs between core and edge switches is possible, but is not a recommended practice.

The long distance Extended Fabric configuration can be established between the following switches:

- HP StorageWorks SAN Switch 2/8 EL
- HP StorageWorks SAN Switch 2/8V
- HP StorageWorks SAN Switch 2/16
- HP StorageWorks SAN Switch 2/16V and SAN Switch 2/16N
- HP StorageWorks SAN Switch 2/32
- HP StorageWorks Core Switch 2/64
- HP StorageWorks SAN Director 2/128

Long-distance ports consume more buffers than regular ISL ports, which means that configuration of a long-distance port could disable the port itself or other ports in the same quad due to lack of buffers.

Note: Long distance between the SAN Switch 2/8 EL, SAN Switch 2/8V, SAN Switch 2/16, SAN Switch 2/16V, SAN Switch 2/16N, SAN Switch 2/32, Core Switch 2/64, and SAN Director 2/128 ports is *not* supported when the long distance fabric-wide parameter fabric.ops.mode.longDistance is set.

For information about configuring an Extended Fabric connection, refer to the *HP StorageWorks Fabric OS 4.2.x Procedures User Guide*.

Note: ISL Trunking is *not* supported on a long-distance ISL.

Remote Switch

6

Remote Switch is an optionally licensed product that enables you to connect two remote HP StorageWorks fabrics over an IP network, enabling communication of IP or ATM protocols as well as fibre channel traffic.

Remote Switch functions with the aid of a “bridging device” or fibre channel gateway. The gateway supports both a fibre channel physical interface and a secondary, non-fibre channel physical interface, such as IP, SONET, or ATM. Remote Switch functions over E_Port connections. With Remote Switch on both fabrics, the gateway accepts fibre channel frames from one fabric, tunnels them across the network, and passes them to the other fabric. From the viewpoint of the connected hosts and storage devices, fabrics using Remote Switch interact the same as locally connected switches.

Remote Switch provides many of the same capabilities of normal ISL links including:

- Coordinated fabric services

 The Remote Switch fabric configuration fully supports all fabric services, including Distributed Name Service, Registered State Change Notification, and Alias Service.

- Distributed management

 Management tools such as Advanced Web Tools, Fabric OS, and SNMP are available from both the local switch and the remote switch. Switch management is routed through the Fibre Channel connection; thus, no additional network connection is required between sites.

- Support for interswitch links (ISLs)

 Sites requiring redundant configurations can connect multiple E_Ports to remote sites by using multiple gateways. Standard Fabric OS routing facilities automatically maximize throughput and provide automatic failover during interruption on the WAN connection.

The Remote Switch feature operates in conjunction with a gateway. The gateway provides an E_Port interface that links to the E_Port. After the link between the two E_Ports has been negotiated, the gateway E_Port moves to passthrough mode and passes fibre channel traffic from the E_Port to the WAN.

The gateway accepts fibre channel frames from one side of a Remote Switch fabric, transfers them across a WAN, and passes them to the other side of the Remote Switch fabric.

Remote Switch is automatically activated when you enable the license key. The only required action is to connect the fabrics through the gateway device and make sure that the `configure` parameters are compatible with the gateway device. Refer to the *HP StorageWorks Fabric OS 4.2.x Procedures User Guide* for additional information.

glossary

This glossary defines terms used in this guide or related to this product and is not a comprehensive glossary of computer terms.

A

ACC
Accept link service reply. The normal reply to an Extended Link Service request (such as FLOGI), indicating that the request has been completed.

address identifier
A 24-bit or 8-bit value used to identify the source or destination of a frame.

AL_PA
Arbitrated loop physical address. A unique 8-bit value assigned during loop initialization to a port in an arbitrated loop. Alternately, “arbitrated loop parameters.”

alias
A logical grouping of elements in a fabric. An alias is a collection of port numbers and connected devices, used to simplify the entry of port numbers and WWNs when creating zones.

arbitrated loop
A shared 100-Mb/sec Fibre Channel transport structured as a loop. Can support up to 126 devices and one fabric attachment.

area number
In Fabric OS v4.x and above, ports on a switch are assigned a logical area number. Port area numbers can be viewed by entering the `switchshow` command. They are used to define the operative port for many Fabric OS commands: for example, area numbers can be used to define the ports within an alias or zone.

ATM
Asynchronous Transfer Mode. A transport used for transmitting data over LANs or WANs that transmit fixed-length units of data. Provides any-to-any connectivity and allows nodes to transmit simultaneously.

authentication
The process of verifying that an entity in a fabric (such as a switch) is what it claims to be.

B

block
As it applies to Fibre Channel technology, upper-level application data that is transferred in a single sequence.

C

CA
Certificate authority. A trusted organization that issues digital certificates.

cascade
Two or more interconnected Fibre Channel switches. HP StorageWorks 2 Gb switches can be cascaded up to 239 switches, with a recommended maximum of seven interswitch links (no path longer than eight switches).

chassis
The metal frame in which the switch and switch components are mounted.

Class 1 service
The class of frame-switching service for a dedicated connection between two communicating ports (also called “connection-oriented service”). Includes acknowledgement of frame delivery or nondelivery.

Class 2 service
A connectionless class of frame-switching service that includes acknowledgement of frame delivery or nondelivery.

Class 3 service
A connectionless class of frame-switching service that does not include acknowledgement of frame delivery or nondelivery. Can be used to provide a multicast connection between the frame originator and recipients, with acknowledgement of frame delivery or nondelivery.

configuration
(1) A set of parameters that can be modified to fine-tune the operation of a switch. Use the `configshow` command to view the current configuration of your switch.

(2) In Zoning, a zoning element that contains a set of zones. The Configuration is the highest-level zoning element and is used to enable or disable a set of zones on the fabric.

congestion
The realization of the potential of oversubscription. A congested link is one on which multiple devices are contending for bandwidth.

CP
Control processor.

CRC
Cyclic redundancy check. A transmission error check that is included in every data frame.

D

D_ID
Destination identifier. A 3-byte field in the frame header, used to indicate the address identifier of the N_Port to which the frame is headed.

datagram
A Class 3 Fibre Channel service that allows data to be sent quickly to devices attached to the fabric, without receipt confirmation.

defined zone configuration
The set of all zone objects defined in the fabric. Can include multiple zone configurations.

deskew
Related to the Trunking feature. The time difference between traffic traveling over each ISL other than the shortest ISL in the group and traffic traveling over that shortest ISL. The deskew number corresponds to nanoseconds divided by 10. The firmware automatically sets the minimum deskew value of the shortest ISL to 15.

digital certificate
An electronic document issued by a CA (certificate authority) to an entity, containing the public key and identity of the entity. Entities in a secure fabric are authenticated based on these certificates.

domain ID
A unique identifier for all switches in a fabric, used in routing frames. Usually automatically assigned by the principal switch but can be assigned manually. The domain ID for an HP StorageWorks switch can be any integer between 1 and 239.

DWDM
Dense wave division multiplexing. Allows more wavelengths to use the same fiber.

E

E_Port
Expansion port. A type of switch port that can be connected to an E_Port on another switch to create an ISL.

ELS
Extended link service. ELSs are sent to the destination N_Port to perform the requested function or service. ELS is a Fibre Channel standard that is sometimes referred to as "Fibre Channel Physical (FC_PH) ELS."

enabled zone configuration
The currently enabled configuration of zones. Only one configuration can be enabled at a time.

error
As it applies to the Fibre Channel industry, a missing or corrupted frame, timeout, loss of synchronization, or loss of signal (link errors).

exchange
The highest-level Fibre Channel mechanism used for communication between N_Ports. Composed of one or more related sequences, it can work in either one or both directions.

F

F_Port
Fabric port. A port that is able to transmit under fabric protocol and interface over links. Can be used to connect an N_Port to a switch.

fabric
A Fibre Channel network containing two or more switches in addition to hosts and devices. Also referred to as a "switched fabric."

Fabric Manager
An HP StorageWorks optionally licensed software application. Fabric Manager is a GUI that allows for fabric-wide administration and management. Switches can be treated as groups, and actions such as firmware downloads can be performed simultaneously.

Fabric Mode
One of two possible modes for an L_Port, in which the L_Port is connected to another port that is not loop capable, using fabric protocol.

fabric name
The unique identifier assigned to a fabric and communicated during login and port discovery.

fabric services
Codes that describe the communication to and from any well-known address.

fabric topology
The arrangement of switches that form a fabric.

Fabric Watch
An HP StorageWorks optionally licensed software application. Fabric Watch can be accessed through either the command line or Advanced Web Tools, and it provides the ability to set thresholds for monitoring fabric conditions.

failover
Describes the Core Switch 2/64 process of one CP passing active status to another CP. A failover is nondisruptive.

fan-in
The ratio of hosts to storage devices; the view of the SAN from the storage port's perspective.

FCP
Fibre Channel Protocol. Mapping of protocols onto the Fibre Channel standard protocols. For example, SCSI FCP maps SCSI-3 onto Fibre Channel.

FCS switch
Relates to the Secure Fabric OS feature. One or more designated switches that store and manage security parameters and configuration data for all switches in the fabric. They also act as a set of backup switches to the primary FCS switch.

FDMI
Fabric-Device Management Interface. FDMI is a database service provided by the fabric for Nx_Ports. The primary use is by HBA devices that register information about themselves and their ports.

FFFFFA
Well-known Fibre Channel address for a management server.

Fibre Channel
Fibre Channel is a protocol used to transmit data between servers, switches, and storage devices. It is a high-speed, serial, bidirectional, topology-independent, multiprotocol, and highly scalable interconnection between computers, peripherals, and networks.

Fibre Channel transport
A protocol service that supports communication between Fibre Channel service providers.

firmware
The basic operating system provided with the hardware.

FL_Port
Fabric loop port. A port that is able to transmit under fabric protocol and also has arbitrated loop capabilities. Can be used to connect an NL_Port to a switch.

flash
Programmable nonvolatile RAM (NVRAM) memory that maintains its contents without power.

FLOGI
Fabric login. The process by which an N_Port determines whether a fabric is present and, if so, exchanges service parameters with it.

frame
The Fibre Channel structure used to transmit data between ports. Consists of a start-of-frame delimiter, header, optional headers, data payload, cyclic redundancy check (CRC), and end-of-frame delimiter. There are two types of frames: link control frames (transmission acknowledgements and so forth) and data frames.

FRU
Field-replaceable unit. A component that can be replaced onsite.

FS
Fibre Channel service. A service that is defined by Fibre Channel standards and exists at a well-known address. For example, the Simple Name Server is a Fibre Channel service.

FSPF
Fabric shortest path first. The routing protocol for Fibre Channel switches.

FTP
File Transfer Protocol.

Fx_Port
A fabric port that can operate as either an F_Port or FL_Port.

G

gateway
Hardware that connects incompatible networks by providing translation for both hardware and software. For example, an ATM gateway can be used to connect a Fibre Channel link to an ATM connection.

GBIC
Gigabit interface converter. A removable serial transceiver module that allows gigabaud physical-level transport for Fibre Channel and gigabit Ethernet.

Gb/sec
Gigabits per second (1,062,500,000 bits/second).

GB/sec
Gigabytes per second (1,062,500,000 bytes/second).

GUI
A graphic user interface, such as Advanced Web Tools and Fabric Manager.

H

HA
High availability. A set of features in HP StorageWorks switches that is designed to provide maximum reliability and nondisruptive replacement of key hardware and software modules.

HBA
Host bus adapter. The interface card between a server or workstation bus and the Fibre Channel network.

header
A Fibre Channel frame has a header and a payload. The header contains control and addressing information associated with the frame.

host
A computer system that provides end users with services like computation and storage access.

hot-swappable
A hot-swappable component can be replaced under power.

HTTP
Hypertext Transfer Protocol. The standard TCP/IP transfer protocol used on the World Wide Web.

I

idle
Continuous transmission of an ordered set over a Fibre Channel link when no data is being transmitted, to keep the link active and maintain bit, byte, and word synchronization.

in-band
Transmission of management protocol over the Fibre Channel.

initiator
A server or workstation on a Fibre Channel network that initiates communications with storage devices.

intermix
Allows any unused bandwidth in a Class 1 connection.

interswitch link
See ISL.

IP
Internet Protocol. The addressing part of TCP.

ISL
Interswitch link. A Fibre Channel link from the E_Port of one switch to the E_Port of another.

ISL oversub-scription ratio
The ratio of the number of free ports (non-ISL) to the number of ISLs on a switch.

J

JBOD
“Just a bunch of disks.” Indicates a number of disks connected in a single chassis to one or more controllers.

K

K28.5
A special 10-bit character used to indicate the beginning of a transmission word that performs Fibre Channel control and signaling functions. The first seven bits of the character are the comma pattern.

key
A string of data (usually a numeric value) shared between two entities and used to control a cryptographic algorithm. Usually selected from a large pool of possible keys to make unauthorized identification of the key difficult.

key pair
In public key cryptography, a pair of keys consisting of an entity's public and private key. The public key can be publicized, but the private key must be kept secret.

L

L_Port
Loop port. A node port (NL_Port) or fabric port (FL_Port) that has arbitrated loop capabilities. An L_Port can be in either Fabric Mode or Loop Mode.

LAN
Local area network. A network in which transmissions typically take place over fewer than 5 kilometers (3.4 miles).

latency
The time required to transmit a frame. Together, latency and bandwidth define the speed and capacity of a link or system.

LIP
Loop initialization primitive. The signal used to begin initialization in a loop. Indicates either loop failure or node resetting.

loop failure
Loss of signal within a loop for any period of time, or loss of synchronization for longer than the timeout value.

loop initialization
The logical procedure used by an L_Port to discover its environment. Can be used to assign AL_PA addresses, detect loop failure, or reset a node.

Loop Mode
One of two possible modes for an L_Port, in which the L_Port is in an arbitrated loop, using loop protocol. An L_Port in Loop Mode can also be in Participating Mode or Nonparticipating Mode.

looplet
A set of devices connected in a loop to a port that is a member of another loop.

LWL
Long wavelength. A type of fiber optic cabling that is based on 1300-mm lasers and supports link speeds of 1.0625 Gb/sec. Can also refer to the type of GBIC or SFP.

M

Mb/sec
Megabits per second.

metric
A relative value assigned to a route to aid in calculating the shortest path (1000 @ 1 Gb/sec, 500 @ 2 Gb/sec).

MIB
Management Information Base. An SNMP structure to help with device management, providing configuration and device information.

MS
Management Server. The Management Server allows a storage area network (SAN) management application to retrieve information and administer the fabric and interconnected elements, such as switches, servers, and storage devices. The MS is located at the Fibre Channel well-known address FFFFFAh.

multicast
The transmission of data from a single source to multiple specified N_Ports (as opposed to all the ports on the network).

N

N_Port
Node port. A port on a node that can connect to a Fibre Channel port or to another N_Port in a point-to-point connection.

Name Server
Simple Name Server (SNS). A switch service that stores names, addresses, and attributes for up to 15 minutes and provides them as required to other devices in the fabric. SNS is defined by Fibre Channel standards and exists at a well-known address. Also referred to as “directory service.”

NL_Port
Node loop port. A node port that has arbitrated loop capabilities. Used to connect an equipment port to the fabric in a loop configuration through an FL_Port.

node
A Fibre Channel device that contains an N_Port or NL_Port.

node name
The unique identifier for a node, communicated during login and port discovery.

Nonparticipating Mode
A mode in which an L_Port in a loop is inactive and cannot arbitrate or send frames but can retransmit received transmissions. This mode is entered if there are more than 127 devices in a loop and an AL_PA cannot be acquired.

NOS
Not operational. The NOS primitive sequence is transmitted to indicate that the FC_Port transmitting the NOS has detected a link failure or is offline, waiting for the offline sequence (OLS) to be received.

NS
Name Server. The service provided by a fabric switch that stores names, addresses, and attributes related to Fibre Channel objects. Can cache information for up to 15 minutes. Also known as “Simple Name Server” or as a “directory service.”

Nx_Port
A node port that can operate as either an N_Port or NL_Port.

O

OLS
Primitive sequence offline.

ordered set
A transmission word that uses 8B/10B mapping and begins with the K28.5 character. Ordered sets occur outside of frames and include the following items:

- **Frame delimiters.** Mark frame boundaries and describe frame contents.
- **Primitive signals.** Indicate events.
- **Primitive sequences.** Indicate or initiate port states.

Ordered sets are used to differentiate Fibre Channel control information from data frames and to manage frame transport.

originator
The Nx_Port that originated an exchange.

out-of-band
Transmission of management protocol outside of the Fibre Channel network, usually over Ethernet.

oversub-scription
A situation in which more nodes could potentially contend for a resource than the resource could simultaneously support (typically an ISL). Oversubscription could be a desirable attribute in fabric topology, as long as it does not produce unacceptable levels of congestion.

P

packet
A set of information transmitted across a network.

parallel
The simultaneous transmission of data bits over multiple lines.

Participating Mode
A mode in which an L_Port in a loop has a valid AL_PA and can arbitrate, send frames, and retransmit received transmissions.

path selection
The selection of a transmission path through the fabric. HP StorageWorks switches use the FSPF protocol.

payload
A Fibre Channel frame has a header and a payload. The payload contains the information being transported by the frame; it is determined by the higher-level service or FC_4 upper-level protocol. There are many different payload formats, based on protocol.

Performance Monitoring
An HP StorageWorks optionally licensed software feature that monitors port traffic and includes frame counters, SCSI read monitors, SCSI write monitors, and other types of monitors.

PLOGI
Port login. The port-to-port login process by which initiators establish sessions with targets.

point-to-point
A Fibre Channel topology that employs direct links between each pair of communicating entities.

port
In an HP StorageWorks switch environment, an SFP or GBIC receptacle on a switch to which an optic cable for another device is attached.

port address
In Fibre Channel technology, the port address is defined in hexadecimal. In the HP StorageWorks Fabric OS, a port address can be defined by a domain and port number combination or by area number. In an ESCON Director, an address used to specify port connectivity parameters and to assign link addresses for attached channels and control units.

port card
A hardware component that provides a platform for field-replaceable, hot-swappable ports.

port log
A record of all activity on a switch, kept in volatile memory.

port name
A user-defined alphanumeric name for a port.

port swapping
Port swapping is the ability to redirect a failed port to another port. This feature is available in Fabric OS v4.1.x and higher.

primary FCS switch
Relates to the HP StorageWorks Secure Fabric OS feature. The primary fabric configuration server switch actively manages security and configurations for all switches in the fabric.

primitive sequence
An ordered set that is transmitted repeatedly and continuously. Primitive sequences are transmitted to indicate specific conditions within or conditions encountered by the receiver logic of an FC_Port.

principal switch
The first switch to boot up in a fabric. Ensures unique domain IDs among roles.

private device
A device that supports arbitrated loop protocol and can interpret 8-bit addresses but cannot log in to the fabric.

private key
The secret half of a key pair.

private loop
An arbitrated loop that does not include a participating FL_Port.

private loop device
A device that supports a loop and can understand 8-bit addresses but does not log in to the fabric.

private NL_Port
An NL_Port that communicates only with other private NL_Ports in the same loop and does not log in to the fabric.

protocol
A defined method and set of standards for communication. Determines the type of error-checking, the data-compression method, how sending devices indicate an end of message, and how receiving devices indicate receipt of a message.

public device
A device that supports arbitrated loop protocol, can interpret 8-bit addresses, and can log in to the fabric.

public key
The public half of a key pair.

public key cryptography
A type of cryptography that uses a key pair, with the two keys in the pair called at different points in the algorithm. The sender uses the recipient's public key to encrypt the message, and the recipient uses the recipient's private key to decrypt it.

public loop
An arbitrated loop that includes a participating FL_Port and can contain both public and private NL_Ports.

Q

QLA
A type of Fibre Channel controller.

QLFA
QuickLoop Fabric Assist. Arbitrated loop technology.

quad
A group of four adjacent ports that share a common pool of frame buffers.

QuickLoop Mode
Allows initiator devices to communicate with private or public devices that are not in the same loop.

R

R_A_TOV
Resource allocation timeout value. The maximum time a frame can be delayed in the fabric and still be delivered.

RAID
Redundant array of independent disks. A collection of disk drives that appear as a single volume to the server and are fault tolerant through mirroring or parity checking.

RCS
Reliable Commit Service. Refers to HP StorageWorks-specific ILS command code.

receiver
A device that performs detection and signal processing.

redundancy
Having multiple occurrences of a component to maintain high availability (HA).

remote switch
An optional product for long-distance fabrics, requiring a Fibre Channel-to-ATM or SONET gateway.

request rate
The rate at which requests arrive at a servicing entity.

route
As it applies to a fabric, the communication path between two switches. Might also apply to the specific path taken by an individual frame, from source to destination. *See also* FSPF.

routing
The assignment of frames to specific switch ports, according to frame destination.

RSCN
Registered state change notification. A switch function that allows notification of fabric changes to be sent from the switch to specified nodes. The fabric controller issues RSCN requests to N_Ports and NL_Ports, but only if they have registered to be notified of state changes in other N_Ports and NL_Ports. This registration is performed via the State Change Registration (SCR) Extended Link Service. An N_Port or NL_Port can issue an RSCN to the fabric controller without having completed SCR with the fabric controller.

S

S_ID
Source ID. Refers to the native port address (24 bit address).

SAN
Storage area network. A network of systems and storage devices that communicate using Fibre Channel protocols.

SCR
State change registration. Extended Link Service (ELS) requests the fabric controller to add the N_Port or NL_Port to the list of N_Ports and NL_Ports registered to receive the Registered State Change Notification (RSCN) Extended Link Service.

SCSI
Small Computer Systems Interface. A parallel bus architecture and a protocol for transmitting large data blocks to a distance of 15 to 25 meters.

SCSI-3
An SCSI standard that defines transmission of SCSI protocol data over different kinds of links.

Secure Fabric OS
A separately sold feature that provides advanced, centralized security for a fabric.

security policy
Rules that determine how security is implemented in a fabric. Security policies can be customized through Secure Fabric OS or Fabric Manager.

sequence
A group of related frames transmitted in the same direction between two N_Ports.

serial
The transmission of data bits in sequential order over a single line.

SES
SCSI Enclosure Services. A subset of the SCSI protocol used to monitor temperature, power, and fan status for enclosed devices.

SFP
Small-form-factor pluggable. A transceiver used on 2 GB/sec switches that replaces the GBIC.

Simple Name Server (SNS)
A switch service that stores names, addresses, and attributes for up to 15 minutes and provides them as required to other devices in the fabric. SNS is defined by Fibre Channel standards and exists at a well-known address. Also referred to as "directory service" or "Name Server."

SMF
Single-mode fiber.

SNMP
Simple Network Management Protocol. An Internet management protocol that uses either IP for network-level functions and UDP for transport-level functions, or TCP/IP for both. Can be made available over other protocols, such as UDP/IP, because it does not rely on the underlying communication protocols.

SNS
Simple Name Server.

soft zone
A zone consisting of zone members that are made visible to each other through client service requests. Typically, soft zones contain zone members that are visible to devices using Name Server exposure of zone members. The fabric does not enforce a soft zone. Note that well-known addresses are implicitly included in every zone.

SONET
Synchronous optical network. A standard for optical networks that provides building blocks and flexible payload mappings.

switch
A fabric device providing bandwidth and high-speed routing of data via link-level addressing.

switch port
A port on a switch. Switch ports can be E_Ports, F_Ports, or FL_Ports.

SWL
Short wavelength. A type of fiber optic cabling that is based on 850 mm lasers and supports
1.0625 GB/sec. link speeds. Can also refer to the type of GBIC or SFP.

T

T11
A standards committee chartered with creating standards for Fibre Channel.

target
A storage device on a Fibre Channel network.

TCP/IP
Transmission Control Protocol / Internet Protocol.

telnet
A virtual terminal emulation used with TCP/IP. “Telnet” is sometimes used as a synonym for the Fabric OS CLI.

throughput
The rate of data flow achieved within a cable, link, or system. Usually measured in bps (bits per second, or b/sec).

Time Server
A Fibre Channel service that allows for the management of all timers.

topology
As it applies to Fibre Channel technology, the configuration of the Fibre Channel network and the resulting communication paths allowed. There are three possible topologies:

transceiver
A device that converts one form of signaling to another for transmission and reception; in fiber optics, optical to electrical.

Translative Mode
A mode in which private devices can communicate with public devices across the fabric.

transmission character
A 10-bit character encoded according to the rules of the 8B/10B algorithm.

transmission word
A group of four transmission characters.

trunking
In Fibre Channel technology, a feature that enables distribution of traffic over the combined bandwidth of up to four ISLs between adjacent switches, while preserving in-order delivery.

trunking group
A set of up to four trunked ISLs.

trunking ports
The ports in a set of trunked ISLs.

tunneling
A technique for enabling two networks to communicate when the source and destination hosts are both on the same type of network but are connected by a different type of network.

U

UDP
User Datagram Protocol. A protocol that runs on top of IP and provides port multiplexing for upper-level protocols.

W

WAN
Wide area network.

well-known address
As it pertains to Fibre Channel technology, a logical address defined by Fibre Channel standards as assigned to a specific function and stored on the switch.

WWN
World Wide Name. An identifier that is unique worldwide. Each entity in a fabric has a separate WWN.

Z

zone
A set of devices and hosts attached to the same fabric and configured as being in the same zone. Devices and hosts within the same zone have access to others in the zone but are not visible to any outside the zone.

zone configuration
A specified set of zones. Enabling a configuration enables all zones in that configuration.

zoning
A feature in fabric switches or hubs that allows segmentation of a node by physical port, name, or address.

index

A

adding end-to-end monitors on a port 51
Advanced Web Tools
 ISL Trunking 74
 overview 21
 performance monitoring 54
 zoning 44
AL_PA performance monitoring 49
API 24
ATM gateway 82
audience 8
authorized reseller, HP 12

C

cfgEnable command 31
cfgSave command 31
cfgTransAbort command 44
commands
 cfgEnable 31
 cfgSave 31
 cfgTransAbort 44
 for ISL Trunking 76
 for managing zoning 45
 for performance monitoring 58
 perfaddusermonitor 53
 portcfglongdistance 65
 portcfgspeed 76
 portperfshow 71
 portzoneshow 34
 urouteconfig 64
conventions
 document 9
 equipment symbols 10
 text symbols 9
core/edge topology and ISL trunking 72
customizing zoning environments 28

D

designing fabric for trunking 72
deskew values 64
disabled zone configuration 32
document
 conventions 9
 related documentation 8
DWDM 77

E

effective zone configuration 31
end-to-end monitors
 adding 51
 setting a mask 52
end-to-end performance monitoring 49
equipment symbols 10
Extended Fabrics 79

F

Fabric Access 24
fabric management recommendations 15
Fabric Manager 22
Fabric Watch 23
fabric, designing for trunking 72
fabric-based zoning 33, 39

FAQs, ISL Trunking 77
filter-based performance monitoring 53
FSPF 64

G

gateway, Remote Switch 82
getting help 12
graphs for performance monitoring 55
groups, trunking 64

H

HA. See high availability
hard zoning 33
hardware-enforced zoning 33
help, obtaining 12
high availability 16
host-based zoning 40
HP
 authorized reseller 12
 storage web site 12
 technical support 12

I

in-order traffic delivery 64
ISL Trunking
 commands 76
 FAQs 77

L

latency 64
link failure 64
long distance ISLs 79
LUN masking 39
LWL
 ISL Trunking support for 62
 trunking with SWL 73

M

Management Server 25
managing ISL Trunking, supported tools 75
managing zones 44
mask for end-to-end monitors 52
master ISL 64
master ports 65
memory constraints for zoning 29
multiple zones 28

N

name server zoning 33, 39

O

optimizing resources through zoning 28
overlapping zones 28

P

perfAddUserMonitor command 53
perfaddusermonitor command 53
performance graphs 55
portcfglongdistance command 65
portcfgspeed command 76
portperfshow command 71
portzoneshow command 34
principal ISL 77

R

rack stability, warning 12
related documentation 8
Remote Switch 81
routing 64
RSCN delivery and zoning 42

S

saved zone configuration 31
security and zoning 43
setting mask for end-to-end monitors 52
slave ISLs 64
slave ports 65
soft zoning 33
software-enforced zoning 33
static routes 64
statistics, port 77
storage-based zoning 39

SWL
ISL Trunking support for 62
trunking with LWL 73
symbols in text 9
symbols on equipment 10

T

technical support, HP 12
text symbols 9
traffic patterns
evaluating 71
planning for 73
traffic, routing of 64
trunking
groups, about 64
ISLs 64
ports 65

U

urouteconfig command 64
user access levels 19
user-defined filter-based monitors 53

W

warning
rack stability 12
symbols on equipment 10
web sites
HP storage 12

Z

zone aliases 31
zone configurations 28, 31
zone enforcement 33
zone management 44
zone objects 28, 30
zones
multiple 28
overlapping 28
zoning
administering security 43
and RSCN delivery 42
customization 28
memory constraints 29
telnet commands 45
zoning schemes 32

Figures

1 Fabric with Three Zones 29
2 Nonoverlapping Hardware Enforced Zoning 35
3 Overlapping Hardware-Enforced Zoning 36
4 Zoning with Hardware Assist (Mixed Port and WWN Zones) 37
5 Overlapping Hardware Enforced Zoning with Soft Porting 37
6 Setting End-to-End Monitors on a Port 51
7 Proper Placement of End-to-End Performance Monitors 52
8 Accessing Performance Graphs 56
9 Eight Performance Monitoring Graphs 57
10 Distribution of Traffic Over an ISL Trunking Group 63
11 Port Groupings on a SAN Switch 2/8 EL 66
12 Quad and Port Groupings on a SAN Switch 2/8V 66
13 Quad and Port Groupings on a SAN Switch 2/16 67
14 Quad and Port Groupings on a SAN Switch 2/16V 67
15 Quad and Port Groupings on a SAN Switch 2/32 68
16 Port Groupings on the Port Card on a Core Switch 2/64 69
17 Port Groupings on the Port Card on a SAN Director 2/128 70
18 Trunk Information Tab in the Switch Admin Window of Advanced Web Tools 75

Tables

1 Document Conventions 9
2 Maximum Allowed Sessions 19
3 RSCN Delivery 42
4 Commands Used for Managing Zoning 45
5 Predefined Performance Graphs 55
6 APM Telnet Commands. 58
7 Fabric OS Commands Related to ISL Trunking 76
8 Frequently Asked Questions about ISL Trunking 77